

VERITAS NetBackup™ 6.0 Advanced Client

System Administrator's Guide

for UNIX, Windows, and Linux

N15272C

September 2005

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 2002 - 2005 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo, and NetBackup are trademarks or registered trademarks of VERITAS Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

Portions of this software are derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm. Copyright 1991-92, RSA Data Security, Inc. Created 1991. All rights reserved.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000
Fax 650-527-2908
www.veritas.com

Third-Party Copyrights

For a list of third-party copyrights, see the *NetBackup Release Notes* appendix.

Contents

Preface	xiii
Getting Help	xiii
Finding NetBackup Documentation	xiii
Accessing the VERITAS Technical Support Web Site	xiii
Contacting VERITAS Licensing	xv
Accessibility Features	xv
Comment on the Documentation	xvi
Advanced Client Assistance	xvi
NDMP Information on the Web	xvi
 Chapter 1. Introduction	1
Overview	2
Features Included in Advanced Client	2
Snapshots	2
Offhost Backup	2
Instant Recovery	4
FlashBackup	4
BLI	4
Advanced Client and NDMP	4
Snapshot Basics	4
Copy-on-Write	5
Mirror	6
Snapshot methods	7
Local Backup of Snapshot	8



Offhost Backup Overview	9
File/Volume Mapping Methods	9
Offhost Backup Methods	10
Alternate Client Backup	10
Data Sharing Between Clients	11
NetBackup Media Server Data Mover (UNIX Only)	15
Third-Party Copy Device Data Mover (UNIX Only)	16
Network Attached Storage Data Mover	17
Offhost Backup Without a SAN (UNIX Only)	17
Offhost Backup with Multi-Ported Array (UNIX Only)	18
Features and Required Software	19
Requirements	21
Restrictions	21
Terminology	23
Chapter 2. Installation	29
Prerequisites	29
Installing Advanced Client On UNIX	30
Loading From Media	30
Distributing Advanced Client Software to UNIX Clients	31
Installing Advanced Client On Windows	33
Distributing Client Software in Mixed-Platform Environments	34
Creating Log Directories	35
Upgrading from Earlier Releases	35
Uninstalling Advanced Client	35
Server Uninstall - UNIX	35
Client Uninstall - UNIX	36
Server/Client Uninstall - Windows	36
Chapter 3. SAN Configuration for Advanced Client	37
SAN Configuration Diagram	38



Supported Peripherals	39
Media Server/Third-Party Copy Requirements	39
Diagram for NetBackup Media Server	40
Diagram for Third-Party Copy Device	41
Diagram for Third-Party Copy Device - Remote	42
Configuration Flowcharts	43
Verify NetBackup Access to SAN Devices	47
Solaris only: Example for sg.links, sg.conf, and st.conf files	50
Device Checklist	52
Solaris only: Configure HBA Drivers	53
Create Backup Configuration Files	54
The 3pc.conf and mover.conf Files: An Overview	54
3pc.conf Description	55
Example 3pc.conf file	55
Determining Requirements	58
What bptpcinfo Automatically Provides	58
What the Backup Methods Require	59
mover.conf Description	59
Types of entries in mover.conf	59
For sites that have one third-party copy device	61
For sites that have multiple third-party copy devices	63
SCSI Reserve/Release	63
Keywords in Mover File	64
Naming the Mover File	68
Selection Priority for mover.conf files	68
Create the 3pc.conf File	69
Create the mover.conf File	72
Chapter 4. Policy Configuration	77
Notes and Prerequisites	78



Configuring an Advanced Client Policy	79
Automatic Snapshot Selection	84
Selecting the Snapshot Method	86
Configuring Backup Scripts	92
Configuring Alternate Client Backup	93
Basic Requirements	93
Available Backup Methods and Snapshot Methods	94
Example configurations	95
Before running the alternate client backup	96
Configuration Tips	96
Maximum Pathname Length	96
Snapshot Tips	96
Backup Policy Configuration Wizard	97
Multiple Data Streams	97
Incremental Backup of Mirror-Based Snapshots	98
Archive Bit Incrementals (Windows only)	98
Access Time Not Updated On Mirror After Backup	98
Chapter 5. FlashBackup Configuration	99
FlashBackup Capabilities	100
Restrictions	100
Configuring a FlashBackup Policy	101
Configuring FlashBackup in the Earlier Manner (UNIX only)	104
Snapshot Selection in Earlier Manner	104
Cache Partition in Earlier Manner	104
Using Multiple Data Streams	106
Chapter 6. NAS Snapshot Configuration	109
NAS Snapshot Overview	110
SnapVault Overview	111
Notes and Prerequisites	112



Setting up a Policy for NAS Snapshots	113
Configuring SnapVault	116
SnapVault Restrictions	116
SnapVault Prerequisites	116
Mount Primary Subvolume On Client	116
Enable NetBackup Access to the SnapVault Primary	117
Enable Access Between SnapVault Primary and SnapVault Secondary	117
Create a Storage Unit for SnapVault	118
Relationships Between Primary and Secondary SubVolumes	120
When does NetBackup perform a full SnapVault?	120
When does NetBackup create a new subvolume for the SnapVault?	120
Notes on SnapVault	121
NAS Snapshot Naming Scheme	122
 Chapter 7. Instant Recovery Configuration	123
Instant Recovery Capabilities	124
Requirements	124
Restrictions	125
Instant Recovery Overview	125
Snapshot and Backup	126
Maintaining the NetBackup Catalogs	126
VxFS Storage Checkpoints	127
VxVM Split Mirrors	127
Snapshot Rotation	127
Configuring Snapshot Deletion	128
Maximum Snapshots Setting (Advanced Snapshot Options dialog)	128
Backup Retention Period	129
Examples for Retention Period and Maximum Snapshots	129
Configuring a Policy for Instant Recovery	131
Configuring VxVM	135



Creating a Snapshot Mirror	135
Creating an Instant Snapshot	136
Creating space-optimized snapshots	136
Creating full-sized snapshots	137
Using the VxVM Graphical User Interface	138
Using VERITAS Enterprise Administrator (VxVM 3.5)	138
Instant Recovery for Databases	140

Chapter 8. Snapshot Configuration Notes 141

nbu_snap	142
Cache device	142
Sizing the Cache Partition	143
How to Enter the Cache	145
VxFS_Checkpoint	147
VxFS Multi-Device System	147
Storage Checkpoint Disk Usage	147
Checkpoint Retention Schedules	148
Block-Level Restore	148
VxFS_Snapshot	149
vxvm	150
Creating a Snapshot Mirror of the Source	150
VxVM Instant Snapshots	151
NetBackup Snapshot Methods	152
Space-Optimized Instant Snapshots	152
FlashSnap	153
Testing Volumes for FlashSnap	153
VVR	156
Set Up Volume Replication	156
Name Registration	156
Primary/Secondary Disk Group and Volume Names	156



Test the Replication Setup	157
NAS_Snapshot	158
VSS_Transportable	159
Prerequisites	159
Notes and Restrictions on VSS_Transportable	159
Array-Related Snapshot Methods	161
Configuration Checklist	161
Overview	162
The Snapshot Methods	162
Client Data Must Be Mirrored	164
Disk Terms	164
Disk Configuration Requirements	165
Access to Disk Arrays	165
Connection to Disk Array: SCSI and Fibre Channel	165
Multiple Connectivity to EMC Array: Common Serial Number mode	167
Configuring Primary and Secondary Disks	167
EMC CLARiiON	167
EMC Symmetrix / DMX	168
Hitachi and HP Arrays	169
Volume Manager Configuration	182
Disk Label	182
Disk Types	182
Disk Group Clones	184
When Secondary Disks are Split and Synched	185
Best Practices	186
NetBackup Access to Arrays	186
Resynchronizing Disks At End of Backup	186
Hardware-Level Disk Restore	186
Volume Manager Disk Groups	186
Volume Manager with Dynamic Multipathing (DMP)	187



Backups Concurrently Accessing Same Disk (no VxVM)	187
Backups Concurrently Accessing VxVM Volumes	188
Concurrent Access to Same VxVM Volume	188
Concurrent Access to Volume Manager Volumes on Same Disks	188
Avoiding Concurrent Access Conflicts	189
Chapter 9. Notes on Media Server and Third-Party Copy	191
Disk Requirements for Media Server/ Third-Party Copy	192
ALL_LOCAL_DRIVES	192
Storage Units	192
Multiplexing	192
Raw Partition Backups	192
Chapter 10. Backup and Restore Procedures	193
Performing a Backup	194
Automatic Backup	194
Manual Backup	194
User-Directed Backup and Archive	194
Performing a Restore	195
Restores from a FlashBackup Backup	195
Restores in Clustered File System (VxFS on UNIX Only)	196
Instant Recovery Restore	196
Instant Recovery: Block-Level Restore (UNIX Clients Only)	197
Instant Recovery: File Promotion (UNIX Clients Only)	197
Instant Recovery: Fast File Resync (Windows Clients Only)	198
Instant Recovery: Snapshot Rollback	199
Configurations for Restore	206
Restoring from a Disk Snapshot	207
Chapter 11. Troubleshooting	213
Gathering Information and Checking Logs	214

Logging Directories for UNIX Platforms	214
Logging Folders for Windows platforms	215
Contacting VERITAS Customer Support	217
Latest Patches and Updates	217
Important Notes	218
Particular Issues	219
Installation Problem	219
FlashBackup and Status code 13	220
Removing a Snapshot	221
Removing a VxVM Volume Clone	226
Appendix A. Processing Background	229
Processing Before and After the Snapshot	230
Quiescing the System	230
Quiescing the Database Application	231
Quiescing the Stack	232
How Copy-on-Write Works	232
Copy-on-write process	233
Backing up a copy-on-write snapshot	234
Appendix B. Managing nbu_snap (Solaris)	237
Cache for nbu_snap	238
Determining Cache Size	238
Terminating nbu_snap	238
nbu_snap Commands	239
snaapon	239
snaplist	239
snapcachelist	241
snapstat	242
snapoff	242



Index 243



Preface

This guide explains how to install, configure, and use VERITAS NetBackup Advanced Client. This guide is intended for the NetBackup system administrator and assumes a thorough working knowledge of UNIX or Windows, and of NetBackup administration.

Getting Help

You can find answers to questions and get help from the NetBackup documentation and from the VERITAS technical support web site.

Finding NetBackup Documentation

A list of the entire NetBackup documentation set appears as an appendix in the *NetBackup Release Notes*. All NetBackup documents are included in PDF format on the NetBackup Documentation CD.

For definitions of NetBackup terms, consult the online glossary.

▼ To access the NetBackup online glossary

1. In the NetBackup Administration Console, click **Help > Help Topics**.
2. Click the **Contents** tab.
3. Click **Glossary of NetBackup Terms**.

Use the scroll function to navigate through the glossary.

Accessing the VERITAS Technical Support Web Site

The address for the VERITAS Technical Support Web site is <http://support.veritas.com>.

The VERITAS Support Web site lets you do any of the following:



- ◆ Obtain updated information about NetBackup, including system requirements, supported platforms, and supported peripherals
- ◆ Contact the VERITAS Technical Support staff and post questions to them
- ◆ Get the latest patches, upgrades, and utilities
- ◆ View the NetBackup Frequently Asked Questions (FAQ) page
- ◆ Search the knowledge base for answers to technical support questions
- ◆ Receive automatic notice of product updates
- ◆ Find out about NetBackup training
- ◆ Read current white papers related to NetBackup

From <http://support.veritas.com>, you can complete various tasks to obtain specific types of support for NetBackup:

1. Subscribe to the VERITAS Email notification service to be informed of software alerts, newly published documentation, Beta programs, and other services.
 - a. From the main <http://support.veritas.com> page, select a product family and a product.
 - b. Under Support Resources, click **Email Notifications**.

Your customer profile ensures you receive the latest VERITAS technical information pertaining to your specific interests.
2. Locate the telephone support directory at <http://support.veritas.com> by clicking the **Phone Support** icon. A page appears that contains VERITAS support numbers from around the world.

Note Telephone support for NetBackup is only available with a valid support contract. To contact VERITAS for technical support, dial the appropriate phone number listed on the Technical Support Guide included in the product box and have your product license information ready for quick navigation to the proper support group.

3. Contact technical support using e-mail.

- a. From the main <http://support.veritas.com> page, click the **E-mail Support** icon.
A wizard guides you to do the following:
 - ◆ Select a language of your preference
 - ◆ Select a product and a platform
 - ◆ Provide additional contact and product information, and your message
 - ◆ Associate your message with an existing technical support case
- b. After providing the required information, click **Send Message**.

Contacting VERITAS Licensing

For license information, you can contact us as follows:

- ◆ Call 1-800-634-4747 and select option 3
- ◆ Fax questions to 1-650-527-0952
- ◆ In the Americas, send e-mail to amercustomercare@veritas.com.
In the Asia and Pacific areas, send email to apaccustomercare@veritas.com.
In all other areas, send email to internationallicense@veritas.com.

Accessibility Features

NetBackup contains features that make the user interface easier to use by people who are visually impaired and by people who have limited dexterity. Accessibility features include:

- ◆ Support for assistive technologies such as screen readers and voice input (Windows servers only)
- ◆ Support for keyboard (mouseless) navigation using accelerator keys and mnemonic keys

For more information, see the *NetBackup Installation Guide*.



Comment on the Documentation

Let us know what you like and dislike about the documentation. Were you able to find the information you needed quickly? Was the information clearly presented? You can report errors and omissions or tell us what you would find useful in future versions of our manuals and online help.

Please include the following information with your comment:

- ◆ The title and product version of the manual on which you are commenting
- ◆ The topic (if relevant) on which you are commenting
- ◆ Your comment
- ◆ Your name

Email your comment to NBDocs@veritas.com.

Please only use this address to comment on product documentation. See “Getting Help” in this preface for information on how to contact Technical Support about our software.

We appreciate your feedback.

Advanced Client Assistance

From the NetBackup Administration Console

For first-time help creating a policy with Advanced Client features, click on the Master Server name at the top of the left pane and click Create a Snapshot Backup Policy.

From the Web (www.support.veritas.com)

For an Advanced Client document containing an up-to-date list of supported operating systems and peripherals, a snapshot/OS compatibility matrix, configuration notes, and a list of NAS vendors supported for the NAS_Snapshot method and for SnapVault, enter “Advanced client configuration” in the search field. The title of the document is: *VERITAS NetBackup Advanced Client Configuration and Compatibility*.

NDMP Information on the Web

The VERITAS support web site has a pdf document on supported NDMP operating systems and NAS vendors. It also contains configuration and troubleshooting help for particular NAS systems. Go to www.support.veritas.com and enter “NAS Appliance” in the search field.

The document’s title is: *NetBackup for NDMP Supported OS and NAS Appliance Information*.

Introduction

1

This chapter describes NetBackup Advanced Client and contains the following topics.

- ◆ [Overview](#)
- ◆ [Snapshot Basics](#)
- ◆ [Local Backup of Snapshot](#)
- ◆ [Offhost Backup Overview](#)
- ◆ [Offhost Backup Methods](#)
- ◆ [Features and Required Software](#)
- ◆ [Requirements](#)
- ◆ [Terminology](#)

Note For help with first-time setup of Advanced Client, see the *NetBackup Advanced Client Quick Start Guide*.



Overview

Advanced Client combines the features of snapshot backup, FlashBackup, BLI Agent, offhost backup, and Instant Recovery. It supports clients on UNIX, Linux, and Windows platforms, on either Fibre Channel networks (SANs) or traditional LANs.

Features Included in Advanced Client

Starting with NetBackup 5.0, Advanced Client combined the features formerly provided by the following NetBackup add-on products: Core Frozen Image Services, Extended Frozen Image Services (Array Integration Option), Offhost and SAN Data Movement Services, FlashBackup, and Persistent Frozen Image.

Snapshots

A snapshot is a point-in-time disk version of the client's data made almost instantaneously. NetBackup backs up the data from the snapshot image, not directly from the client's primary data. This allows client operations and user access to continue without interruption during the backup.

A snapshot is required for the other features of Advanced Client. A number of methods are provided for creating snapshots. You can select the snapshot method manually from the Policy dialog of the NetBackup Administration Console, or allow NetBackup to select the method for you.

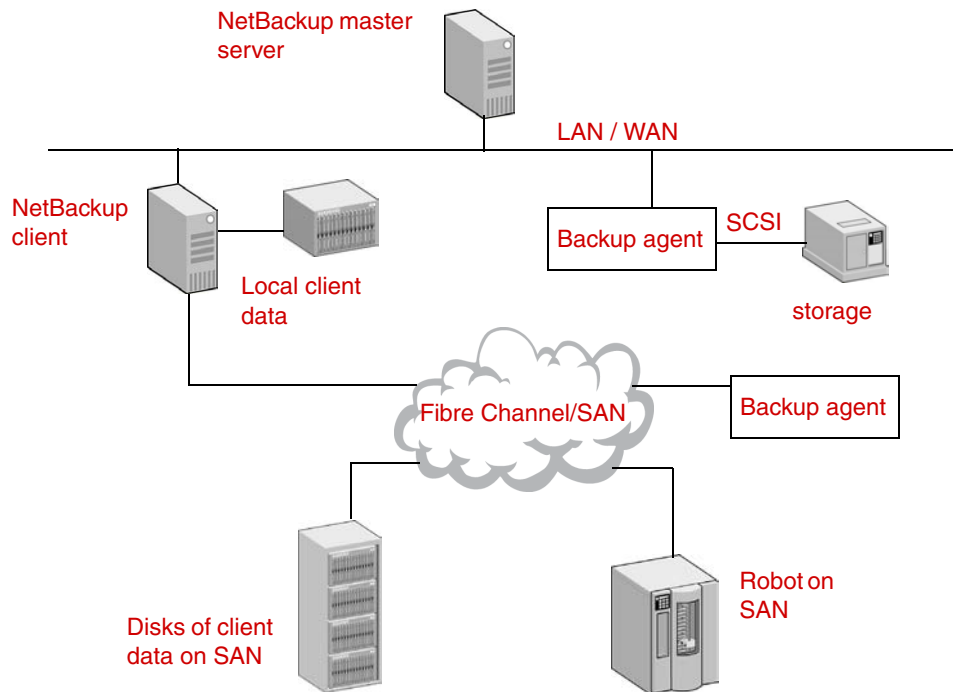
Refer to "[Snapshot Basics](#)" on page 4 for background on snapshot technology.

Offhost Backup

Another major component of NetBackup Advanced Client is support for offhost backup. Offhost backup shifts the burden of backup processing onto a separate *backup agent*, greatly reducing the impact on the client's computing resources ordinarily caused by a local backup. The client supplies a relatively small amount of mapping information, but the backup agent does the bulk of the work by sending the client's actual data to the storage device.

See the following network diagram showing a backup agent.

Backup Agent for Offhost Backup



The backup agent can be any of the following:

- ◆ an additional (alternate) client
- ◆ a NetBackup media server or a third-party copy device that implements the SCSI Extended Copy command.
- ◆ a NAS host (Network Attached Storage)

Note that many types of devices are designed to act as third-party copy devices, such as routers, bridges, robotic libraries, and disk arrays. The backup agent can direct the data to SCSI-attached storage or to storage on the SAN.



Instant Recovery

This feature makes backups available for quick recovery from disk. Instant Recovery combines snapshot technology—the image is created without interrupting user access to data—with the ability to do rapid snapshot-based restores. The image is retained on disk and can also be backed up to tape or other storage.

Instant Recovery makes possible three additional variations of restore: block-level restore, file promotion, and snapshot rollback. For a description, refer to “[Instant Recovery Restore](#)” on page 196.

FlashBackup

FlashBackup is a policy type that combines the speed of raw-partition backups with the ability to restore individual files.

BLI

Block level incremental backup extends the capabilities of NetBackup to back up only changed data blocks of Oracle and DB2 database files. Refer to the appropriate NetBackup database agent guide for details.

Advanced Client and NDMP

Using the NDMP V4 snapshot extension, NetBackup Advanced Client can make policy-based snapshots of data on a Network Attached Storage (NAS) host. The snapshot is stored on the same NAS device that contains the primary client data. From the snapshot, you can restore individual files or roll back an entire volume or file system, by means of Instant Recovery.

Note NetBackup for NDMP add-on software is required, and the NAS vendor must support snapshots.

Snapshot Basics

Large active databases or file systems that must be available around-the-clock are difficult to back up without incurring a penalty. Often, the penalty takes one of two forms:

- ◆ The entire database is taken offline or the file system is unmounted, to allow time for the backup, resulting in suspension of service and inconvenience to users.



- ◆ The copy is made very quickly but produces an incomplete version of the data, some transactions having failed to complete.

A solution to this problem is to create a *snapshot* of the data. This means capturing the data at a particular instant, without causing significant client downtime. The resulting capture or snapshot can be backed up without affecting the performance or availability of the file system or database. Without a complete, up-to-date snapshot of the data, a correct backup cannot be made.

When a backup is managed by a NetBackup media server or third-party copy device on a Fibre Channel network, the data to back up must be contained in a snapshot. The backup agent can only access the data by means of the raw physical disk. Once the data is captured as a snapshot, the NetBackup client “maps” the logical representation of the data to its absolute physical disk address. These disk addresses are sent to the backup agent over the LAN and the data is then read from the appropriate disk by the backup agent. (This process is explained in greater detail under “[Offhost Backup Overview](#)” later in this chapter.)

Two types of snapshots are available, both supported by NetBackup: copy-on-write, and mirror.

Copy-on-Write

A copy-on-write type of snapshot is a detailed account of data as it existed at a certain moment. A copy-on-write is not really a copy of the data, but a specialized “record” of it.

A copy-on-write snapshot is created in available space in the client’s file system or in a designated raw partition, not as a complete copy of the client data on a separate or mirror disk. The snapshot is then backed up to storage as specified in the backup policy. Users can access their data without interruption, as though no backup is taking place. The file system is paused just long enough to assemble a transactionally consistent record.

For a description of the copy-on-write process, see “[How Copy-on-Write Works](#)” on page 232. Note that in VxFS, there are two kinds of copy-on-write snapshots: file system snapshots and Storage Checkpoints.

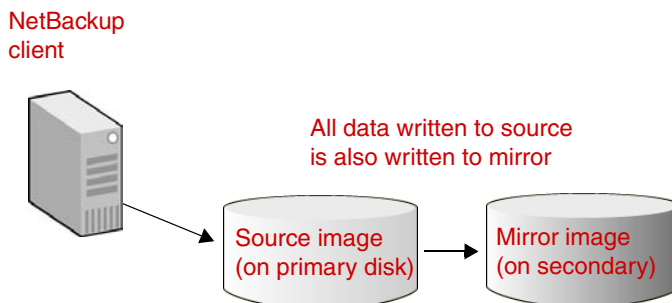
Benefits of copy-on-write:

- ◆ Consumes less disk space: no need for secondary disks containing complete copies of source data.
- ◆ Relatively easy to configure (no need to set up mirror disks).
- ◆ Creates a snapshot much faster than one created by a large, unsynchronized mirror, because mirror synchronization can be time consuming.

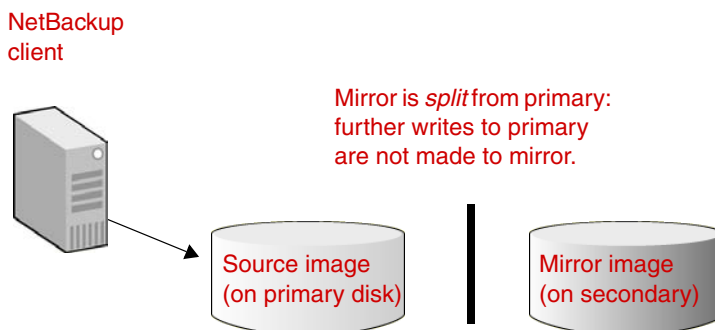


Mirror

Unlike a copy-on-write, a mirror is a complete data copy stored on a separate disk, physically independent of the original. Every change or write to the data on the primary disk is also made to the copy on the secondary disk. This creates a “mirror” image of the source data.



As in a copy-on-write, transactions are allowed to finish and new I/O on the primary disk is briefly halted. When the mirror image is brought up-to-date with the source (made identical to it), the mirror is *split* from the primary, meaning that new changes can be made to the primary but not to the mirror. At this point the mirror can be backed up (see next diagram).



If the mirror will be used again it must be brought up-to-date with the primary volume (resynchronized). During resynchronization, the changes made to the primary volume—while the mirror was split—are written to the mirror.

Since mirroring requires an exact, complete copy of the primary on a separate device (equal in size to the disk being mirrored), it consumes more disk space than a copy-on-write.

Benefits of mirror:

- ◆ Has less impact on the performance of the application or database host being backed up (NetBackup client), because there is no need to run the copy-on-write mechanism.
- ◆ Allows faster backups: the backup process reads data from a separate disk (mirror) operating independently of the primary disk that holds the client's source data. This means that, unlike the copy-on-write, there is no need to share disk I/O with other processes or applications. Apart from NetBackup, no other applications have access to the mirror disk. During a copy-on-write, the source data can be accessed by other applications as well as by the copy-on-write mechanism.

Note If additional disk drives are available and volumes have been configured with the VERITAS Volume Manager, a mirror snapshot method is usually a good choice.

Snapshot methods


Advanced Client supports a number of methods for creating a snapshot image. You can select the method or let NetBackup select it based on your environment. The snapshot methods are described in the [“Policy Configuration”](#) chapter.



Local Backup of Snapshot

A snapshot can be backed up to any NetBackup storage device. A fibre channel network or SAN is not required. The following diagram shows a network configuration sufficient for a local backup of a snapshot. The network configuration is identical to that for normal NetBackup (no snapshot).

Snapshot Backup on Local Network (No Fibre channel/SAN Required)



1. NetBackup master server tells the client to create the snapshot data on the disk.
2. Client sends the data to the media server.
3. Media server processes the backup and reads the client data.
4. Media server writes data to local storage.

Offhost Backup Overview

One of the principal goals of NetBackup Advanced Client is to move I/O processing off the primary NetBackup client to a backup agent. Advanced Client supports several methods of doing this:

- ◆ **Alternate client backup:** a second or “alternate” client performs the backup on behalf of the primary client. Compared to the other offhost methods, this approach reduces the backup I/O burden on the primary client to the greatest extent.
- ◆ **Data mover: NetBackup Media Server** (UNIX clients only): a NetBackup media server reads the backup data from the client snapshot and writes the data to a storage device, using mapping information provided by the client.
- ◆ **Data mover: Third-Party Copy Device** data mover (UNIX clients only): using the Extended Copy command and mapping information from the client, a third-party copy device reads the backup data from the client snapshot and writes the data to a storage device. Many kinds of devices, such as routers and disk arrays, are designed as third-party copy devices. For a list of supported third-party copy devices, refer to “[Advanced Client Assistance](#)” on page xvi.
- ◆ **Data mover: Network Attached Storage:** an NDMP (NAS) host performs the backup.

File/Volume Mapping Methods

If the agent performing the offhost backup is either NetBackup Media Server or Third-Party Copy Device, the agent is unaware of logical organizations of data such as file systems and volume managers, and can access the data only from the physical disk address location. In order for NetBackup to perform this type of backup, it must translate the logical representation of the data to its physical disk addresses. This logical-to-physical translation process is referred to as *mapping* the data. During the backup, the mapping information is transmitted to the media server.

The mapping methods are installed as part of the NetBackup Advanced Client product. When a backup is initiated, NetBackup automatically selects the correct mapping method, depending on whether the backup data is configured over physical devices, logical volumes, or file systems.



Offhost Backup Methods

NetBackup Advanced Client supports several forms of offhost backup, as explained in the following sections.

Alternate Client Backup

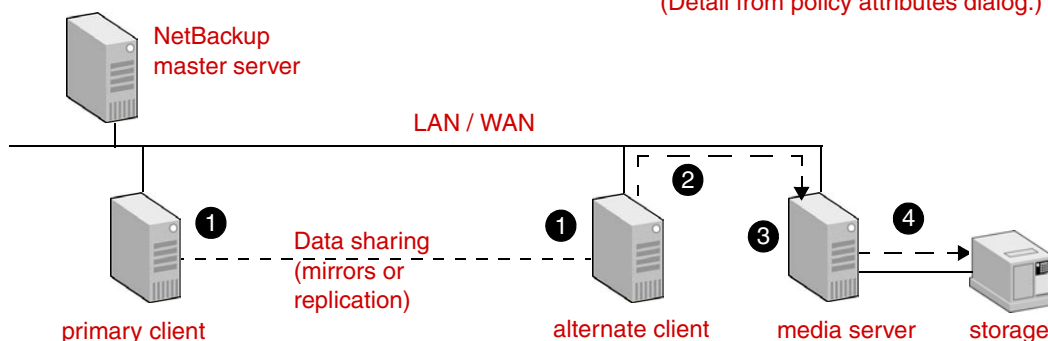
With this feature, all backup processing is off-loaded to another client. Offloading the work to an alternate client saves computing resources on the primary client. The backup I/O processing is handled by the alternate client, and the backup has little or no impact on the primary client.

The following diagram is a basic view of alternate client backup. A NetBackup master server is connected by means of a local or wide-area network to two clients and a media server. The primary or original NetBackup client contains the data to be backed up, and the alternate NetBackup client (which could be at a remote site) has a copy of that data. The media server can be accessed by the alternate client. This means the media server can back up the alternate client locally.

Alternate Client Backup: Backup is performed on alternate client



(Detail from policy attributes dialog.)



1. Primary or alternate client creates the snapshot data on disk.
2. Alternate client sends the snapshot data to the media server.
3. Media server reads the snapshot data from the alternate client.
4. Media server writes data to local storage.

Data Sharing Between Clients

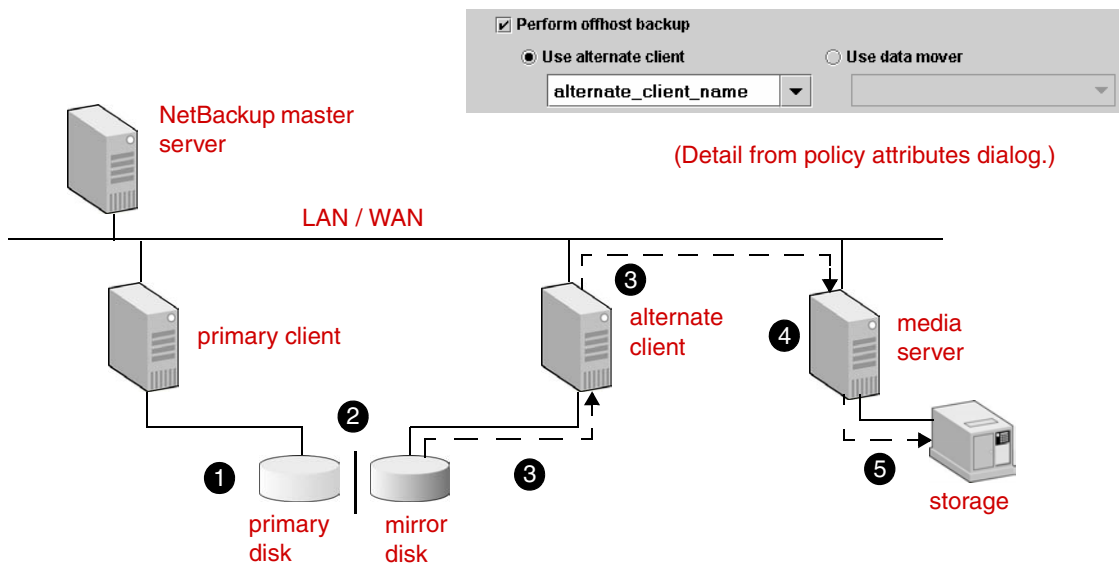
For alternate client backup, the original (primary) and alternate client must share data. Two types of sharing are supported: split mirror, and data replication.

Split mirror

In this configuration, the alternate client has access to mirror disks containing the primary client's data. Before the backup, the mirror is split from the primary disk and the snapshot is made on the mirror disk. The alternate client has access to the mirror disk, and the media server backs up the alternate client. After the backup, the mirror can be optionally resynchronized with the primary disk.

Note The mirror disk need not be visible to the primary client, only to the alternate client.

Alternate Client and Split Mirror: Primary client and alternate client share data through mirroring.



1. Mirror disk is synced with primary.
2. Primary client creates the snapshot by splitting the mirror disk from primary disk.
3. Alternate client sends the snapshot data from the mirror to the media server.
4. Media server reads the snapshot data from the alternate client.
5. Media server writes data to storage.

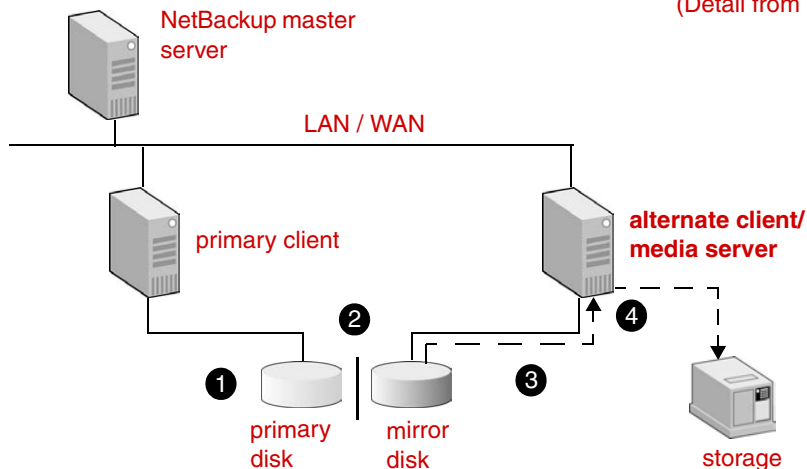


Note that the media server and alternate client can be the same host:

Alternate Client and Media Server on Same Host



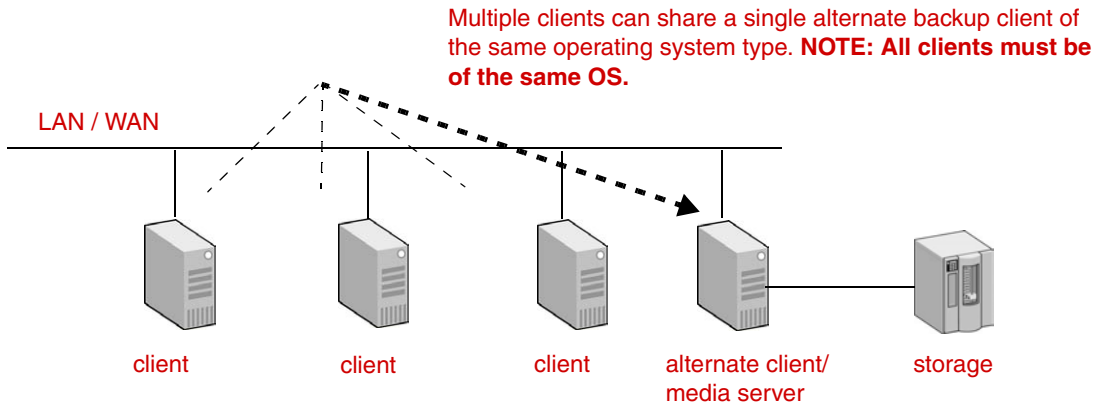
(Detail from policy attributes dialog.)



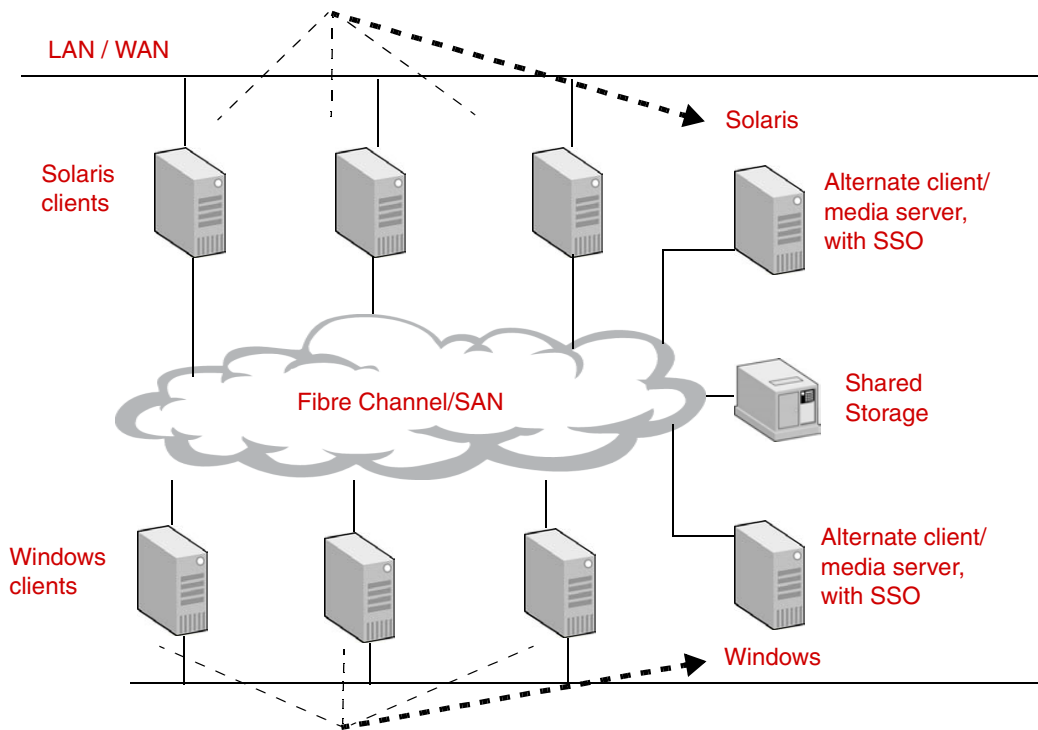
1. Mirror disk is synced with primary.
2. Primary client creates the snapshot by splitting the mirror disk from primary disk.
3. Media server (serving as alternate client) reads the snapshot data from the mirror.
4. Media server writes data to storage.

A single alternate client can handle backups for a number of primary clients, as shown in the following diagram.

Alternate Client for Multiple Primary Clients



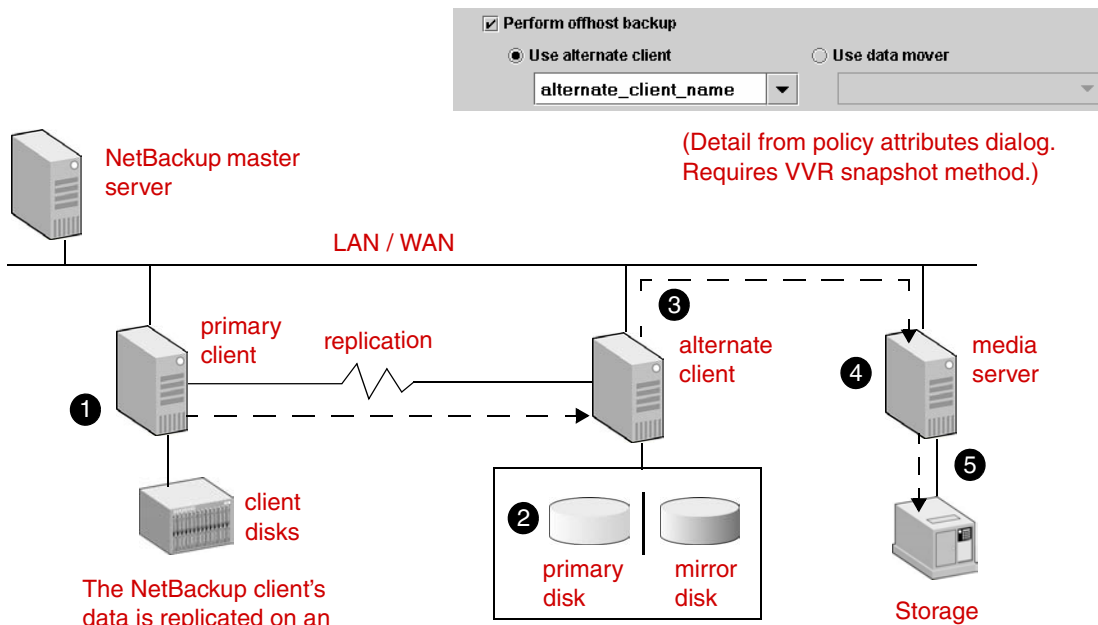
Multiple clients with SSO: Alternate client performs backup for multiple primary clients with NetBackup SSO option on a SAN.



Data Replication (UNIX Only)

In this configuration, a disk on the alternate client keeps a replicated copy of the client's data. When the backup starts, the mirror on the alternate client is split from its primary disk, and the snapshot is made on the mirror. At this point, the media server can back up the alternate client. After the backup, the volume and file system on the mirror are unmounted and the mirror is resynchronized with its primary disk.

Replication: Primary client and alternate client share data through replication



1. Primary client tells alternate client when to create snapshot.
2. Alternate client creates the snapshot by splitting mirror from primary.
3. Alternate client sends the snapshot data from the mirror to the media server.
4. Media server reads the snapshot data from the alternate client.
5. Media server writes data to storage.

The above configuration is supported by the VVR snapshot method for UNIX clients only, and requires the VERITAS Volume Manager (VxVM version 3.2 or later) with the VVR license.

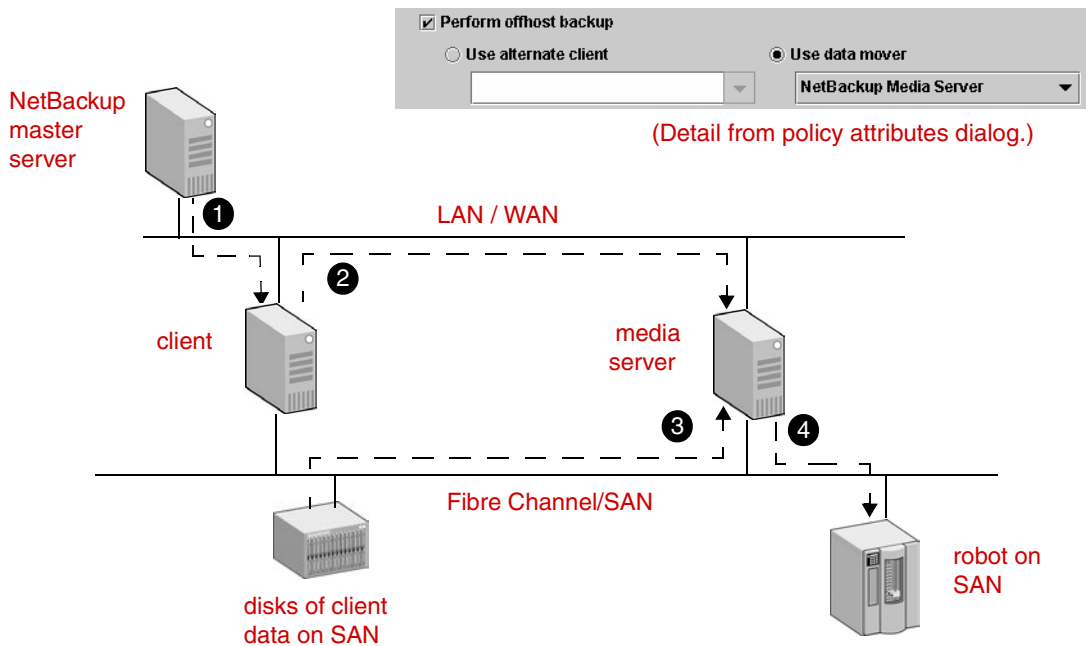
NetBackup Media Server Data Mover (UNIX Only)

In this offhost backup method, a NetBackup media server reads the backup data from the client snapshot and writes the data to a storage device, using mapping information provided by the client. This method is not supported for Windows clients.

This mapping-based offhost data movement is in addition to the normal backup processing done by the media server in a typical master server/media server configuration.

Note If you have a multi-ported SCSI disk array, a fibre channel SAN is not required. See “Offhost Backup Without a SAN (UNIX Only)” on page 17.

NetBackup Media Server



1. On LAN, NetBackup master server tells the client to map the snapshot data on the disk.
2. On LAN, client sends the mapping information to the media server.
3. Media server processes the mapping information and reads client data over the SAN, from the addresses specified by the client.
4. Media server writes data across the SAN to storage.

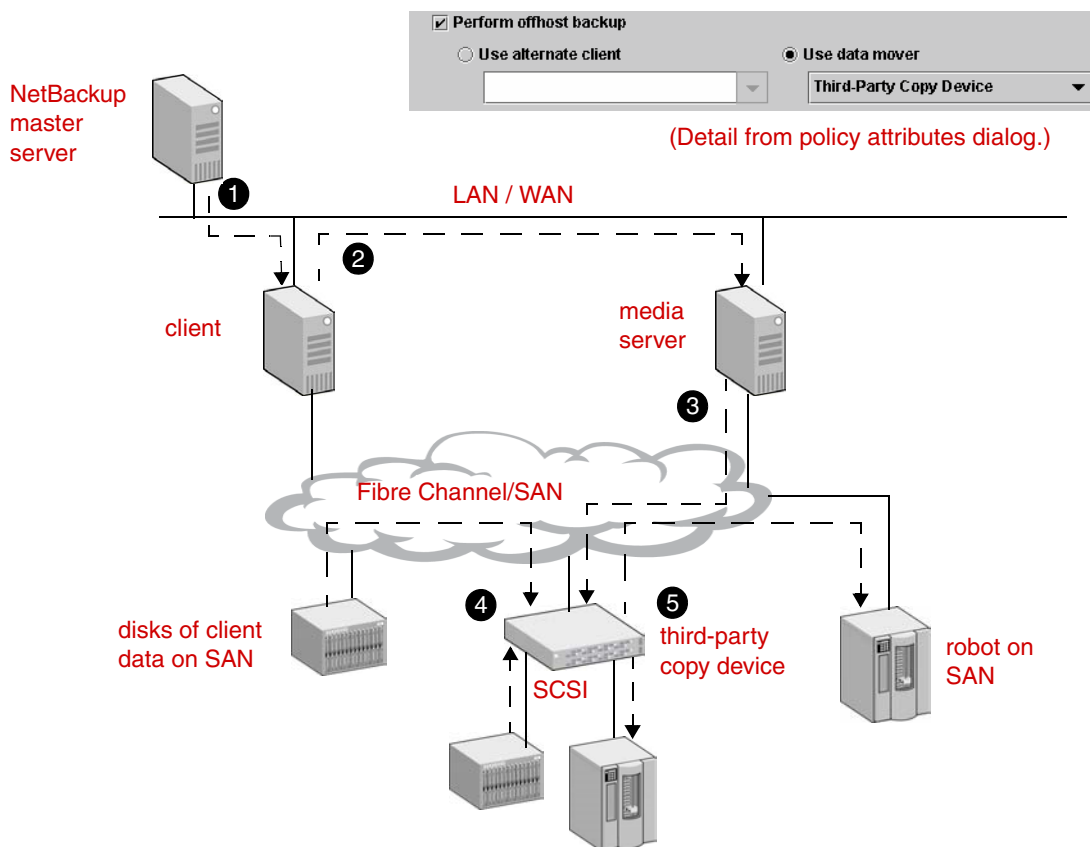


Third-Party Copy Device Data Mover (UNIX Only)

In this backup method, using the Extended Copy command and mapping information from the client, a third-party copy device reads the backup data from the client snapshot and writes the data to a storage device. This method is not supported for Windows clients.

Choose this option if processing time on the NetBackup client is critical and off-loading backup processing to a third-party copy device may save time.

Third-Party Copy



1. On LAN, NetBackup master server tells the client to map the snapshot data on the disk.
2. On LAN, client sends the mapping information to the media server.
3. Media server sends third-party copy commands to the third-party copy device over the SAN.
4. Third-party copy device reads the client data from either SAN-attached or SCSI-attached disk.
5. Third-party copy device writes data to SAN-attached or SCSI-attached storage.

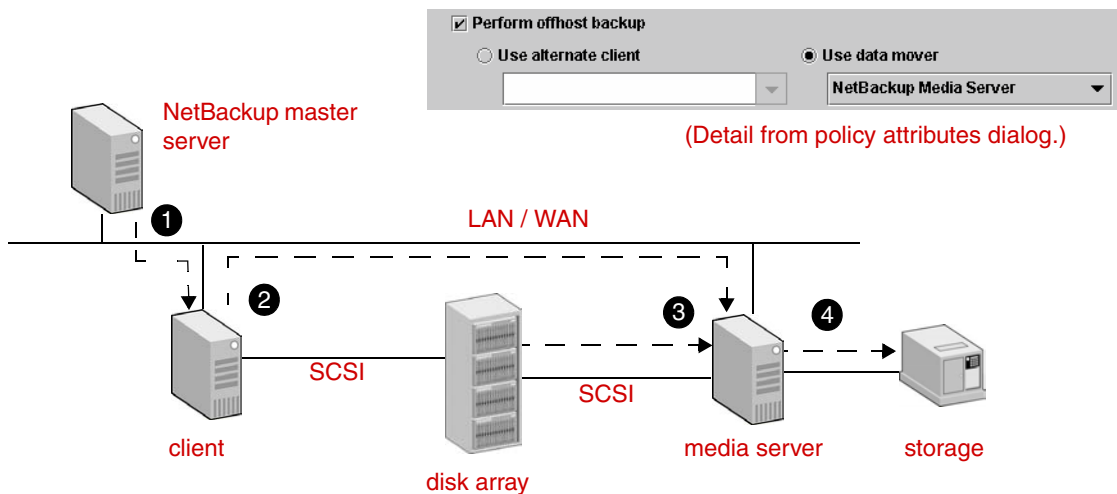
Network Attached Storage Data Mover

In this offhost backup method, an NDMP (NAS) host makes a snapshot of client data and stores the snapshot on the same NAS-attached device that contains the original (primary) data. The snapshot can be restored by means of Instant Recovery.

Offhost Backup Without a SAN (UNIX Only)

Offhost backup does not require a fibre channel SAN. You can configure a multi-ported SCSI disk array on a LAN or WAN (as shown below) to support a NetBackup Media Server offhost backup. The NetBackup media server performs the data movement. This approach is not supported for Windows clients.

NetBackup Media Server with Multi-Ported Disk Array (No SAN)



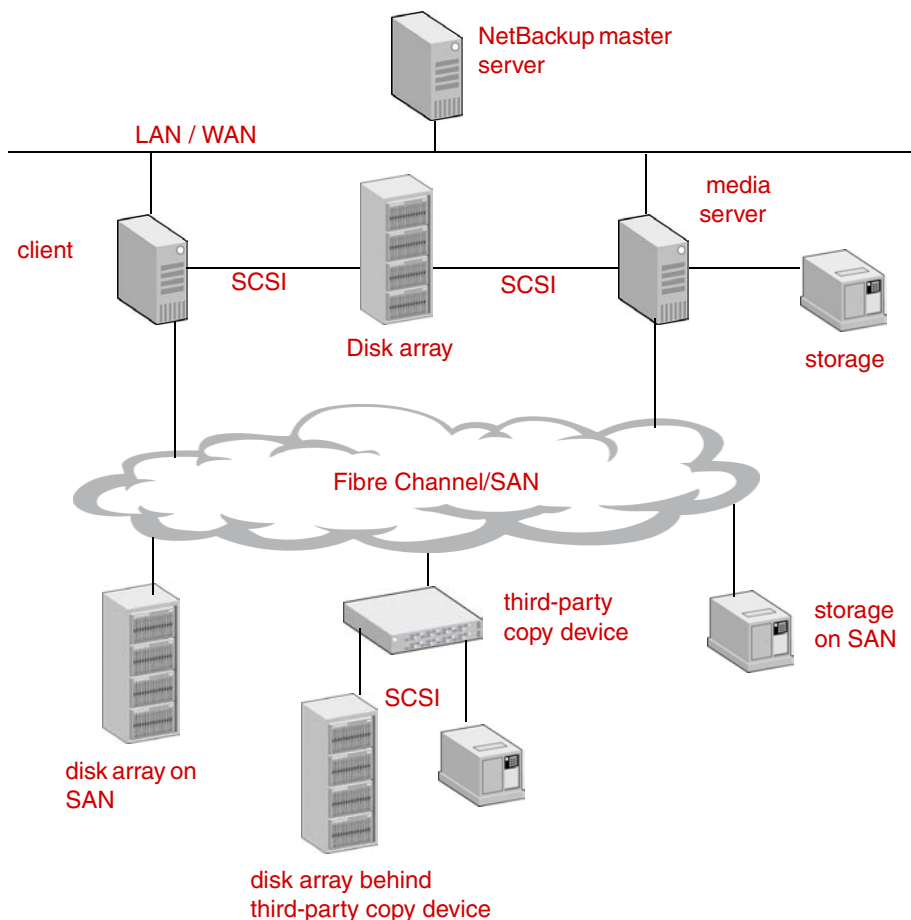
1. NetBackup master server tells the client to map the snapshot data on the disk.
2. Client sends the mapping information to the media server.
3. Media server processes the backup and reads client data from the addresses specified by the client.
4. Media server writes data to storage.



Offhost Backup with Multi-Ported Array (UNIX Only)

A multi-ported disk array can be combined with fibre channel to support either a NetBackup Media Server or Third-Party Copy Device backup. The diagram below shows the following placement options for the disk array:

- ◆ On the LAN using SCSI connections to the NetBackup client and media server
- ◆ On the SAN behind a third-party copy device
- ◆ Directly attached to the SAN



Features and Required Software

The following table shows the types of backup you can configure with Advanced Client, and the corresponding NetBackup features and requirements.

Note Advanced Client supports these policy types: DB2, FlashBackup, FlashBackup-Windows, MS-Exchange, MS-SQL-Server, MS-Windows-NT, Oracle, and Standard.

Advanced Client Features and Requirements

Type of Backup	NetBackup Advanced Client Feature	NetBackup Software Required	Other Software Required
Client and storage devices on local network, and snapshot not required.	<i>Advanced Client not required.</i>	<ul style="list-style-type: none"> NetBackup 6.0. 	
Windows client and storage devices on local network, open file backup required.	<ul style="list-style-type: none"> <i>Advanced Client not required.</i> Snapshot method: VSS, VSP (included in base NetBackup). 	<ul style="list-style-type: none"> NetBackup 6.0. 	
Client and storage devices on local network, snapshot required. Wider selection of snapshots needed for UNIX or Windows clients.	<ul style="list-style-type: none"> Snapshot method: nbu_snap, VxFS_Snapshot, VxFS_Checkpoint, vxvm. For particular arrays: TimeFinder, ShadowImage, BusinessCopy. Optional policy attribute: Retain snapshots for instant recovery. 	<ul style="list-style-type: none"> Advanced Client Additional NetBackup database software needed if backing up clients in a database. 	<ul style="list-style-type: none"> Requirements depend on the snapshot method (see the <i>NetBackup Advanced Client Configuration and Compatibility</i> document on the VERTIAS support site^a). For Oracle clients, Oracle8i or later must be installed.
Local client shares data with alternate client (split-mirror or replication).	<ul style="list-style-type: none"> Snapshot method: FlashSnap, VVR, VSS_Transportable, TimeFinder, ShadowImage, BusinessCopy. Optional policy attribute: Retain snapshots for instant recovery. 	<ul style="list-style-type: none"> Advanced Client Additional NetBackup database software needed if backing up clients in a database. 	<ul style="list-style-type: none"> Requirements depend on the snapshot method (see the <i>NetBackup Advanced Client Configuration and Compatibility</i> document on the VERTIAS support site^a). For Oracle clients, Oracle8i or later must be installed.



Advanced Client Features and Requirements

Type of Backup	NetBackup Advanced Client Feature	NetBackup Software Required	Other Software Required
Client and storage devices on SAN; third-party copy device available. NOTE: Not supported on Windows clients.	<ul style="list-style-type: none"> Snapshot method: nbu_snap, VxFS_Checkpoint, vxvm. For particular arrays: TimeFinder, ShadowImage, BusinessCopy. Data mover: Third-Party Copy Device. Optional policy attribute: Retain snapshots for instant recovery. 	<ul style="list-style-type: none"> Advanced Client Additional NetBackup database software needed if backing up clients in a database. 	<ul style="list-style-type: none"> Requirements depend on the snapshot method (see the <i>NetBackup Advanced Client Configuration and Compatibility</i> document on the VERTIAS support site^a). For Oracle clients, Oracle8i or later must be installed.
Client and storage devices on SAN; third-party copy device <i>not</i> available. NOTE: Not supported on Windows clients.	<ul style="list-style-type: none"> Snapshot method: nbu_snap, VxFS_Checkpoint, vxvm. For particular arrays: TimeFinder, ShadowImage, BusinessCopy. Data mover: NetBackup Media Server. Optional policy attribute: Retain snapshots for instant recovery. 	<ul style="list-style-type: none"> Advanced Client Additional NetBackup database software needed if backing up clients in a database. 	<ul style="list-style-type: none"> Requirements depend on the snapshot method (see the <i>NetBackup Advanced Client Configuration and Compatibility</i> document on the VERTIAS support site^a). For Oracle clients, Oracle8i or later must be installed.
Client data stored on NDMP (NAS) host	<ul style="list-style-type: none"> Snapshot method: NAS_Snapshot. Data mover: Network Attached Storage. Additional required policy attributes: Retain snapshots for instant recovery; and Snapshots only in Schedule. 	<ul style="list-style-type: none"> Advanced Client NetBackup for NDMP 	<ul style="list-style-type: none"> See the <i>NetBackup Advanced Client Configuration and Compatibility</i> document on the VERTIAS support site^a.
Client data stored on NDMP (NAS) host, to be backed up to SnapVault secondary system	<ul style="list-style-type: none"> Snapshot method: NAS_Snapshot. Data mover: Network Attached Storage. Additional required policy attributes: Retain snapshots for instant recovery; and Snapshots only in Schedule. SnapVault disk storage unit. 	<ul style="list-style-type: none"> Advanced Client NetBackup for NDMP 	<ul style="list-style-type: none"> SnapVault primary and secondary hosts must be licensed by NAS vendor. For details, see the “NAS Snapshot Configuration” chapter and the <i>NetBackup Advanced Client Configuration and Compatibility</i> document on the VERTIAS support site^a.

a. For help accessing this document, refer to “[Advanced Client Assistance](#)” on page xvi.

Requirements

NetBackup Advanced Client requires the following components:

- ◆ A master server with NetBackup Advanced Client server software installed.
- ◆ Clients running Solaris, HP, AIX, Linux, or Windows, with NetBackup Advanced Client software installed. For AIX and Linux clients, the client data must be reside in a VxFS file system.

Note Certain operating system and device patches (such as for the host bus adapter) may be required for both servers and clients. To obtain the latest information, refer to “[Advanced Client Assistance](#)” on page xvi.

Please note the following additional requirements:

- ◆ For the VxFS_Checkpoint snapshot method, all clients must have VxFS 3.4 or later with the Storage Checkpoints feature.
- ◆ To use Advanced Client to back up a VxFS file system, the VxFS file system on the client has to be patched with the correct dynamic linked libraries.
- ◆ For the vxvm snapshot method, all clients must have VxVM 3.1 or later.
- ◆ For the FlashSnap and VVR snapshot methods, all clients must have VxVM 3.2 or later. Each method requires its own add-on license to VxVM.
- ◆ For the TimeFinder, ShadowImage, or BusinessCopy snapshot methods, assistance may be required from the disk array vendor. Refer to “[Array-Related Snapshot Methods](#)” on page 161.
- ◆ To use the snapshot and offhost backup features of NetBackup Advanced Client with a NetBackup Oracle policy, UNIX clients must have Oracle8i or later installed.
- ◆ To use the snapshot and offhost backup feature of NetBackup Advanced Client for HP clients, the HP client must be using the Online JFS file system, not the default JFS.

Restrictions

For a complete list of supported peripherals, and for other operational notes, refer to the *NetBackup Release Notes*, or to “[Advanced Client Assistance](#)” on page xvi.

Note the following restrictions:

- ◆ Advanced Client does not support the ALL_LOCAL_DRIVES entry in the policy’s Backup Selections list.
- ◆ VxFS multi-device file systems are only supported by the VxFS_Checkpoint and vxvm snapshot methods.



- ◆ For NetBackup Media Server or Third-Party Copy Device backup method, the disk containing the client's data (the files to back up) must be either a SCSI or Fibre Channel device.
- ◆ For offhost backup using a data mover with the `nbu_snap`, `VxFS_Checkpoint`, or `vxvm` snapshot methods, the NetBackup media server must be able to access all the disks that make up the snapshot image. The disk(s) can be connected to a SAN. For each of these snapshot methods, note the following:
 - ◆ **nbu_snap**: media server requires access to the active disk and the cache disk.
 - ◆ **VxFS_Checkpoint**: media server requires access to the primary or active disk.
 - ◆ **vxvm**: access requirements depend on layout of the volume group. Media server must be able to access all disks that make up the snap mirror volume.
- ◆ For the TimeFinder, ShadowImage, or BusinessCopy snapshot methods (when using the NetBackup Media Server or Third-Party Copy Device backup methods), the NetBackup clients must be able to access the mirror (secondary) disk containing the snapshot of the client's data. The NetBackup clients must also be able to access the primary disk. The NetBackup media server only needs access to the mirror (secondary) disk.
- ◆ For the TimeFinder, ShadowImage, or BusinessCopy snapshot methods, a Volume Manager disk group must consist of disks that are all made by the same vendor.
- ◆ For the NetBackup Media Server or Third-Party Copy Device backup method, the disk must be able to return its SCSI serial number in response to a serial-number inquiry (serialization), or the disk must support SCSI Inquiry Page Code 83.
- ◆ Multiplexing is not supported for Third-Party Copy Device offhost backups.
- ◆ For alternate client backup, the user and group identification numbers (UIDs and GIDs) for the files to be backed up must be available to both hosts (the primary client and the alternate backup client).
- ◆ Inline Tape Copies (called Multiple Copies in Vault) is not supported for Third-Party Copy Device offhost backups.
- ◆ For media servers running AIX (4.3.3 and higher), note the following:
 - ◆ Clients must be Solaris, HP, or AIX.
 - ◆ Requires the use of tape or disk LUNs to send the Extended Copy commands for backup.
 - ◆ The tape must be behind a third-party-copy-capable FC-to-SCSI router, and the router must be able to intercept Extended Copy commands sent to the tape LUNs.
 - ◆ The mover.conf file must have a tape path defined, not a controller path.

Terminology

This section introduces terms used with NetBackup Advanced Client. For explanations of other NetBackup terms, consult the NetBackup online glossary.

Alternate Client Backup

A backup performed by one client on behalf of another client.

Backup agent (see also Third-Party Copy Device)

A general term for the host that manages the backup on behalf of the NetBackup client. This is either another client, the NetBackup media server, a third-party copy device, or a NAS filer.

BCV

The mirror disk in an EMC primary-mirror array configuration (see *mirror*). BCV stands for “Business Continuation Volume.”

Bridge

In a SAN network, a *bridge* connects SCSI devices to Fibre Channel. A *third-party copy device* can be implemented as part of a bridge or as part of other devices. Note that not all bridges function as third-party copy devices.

BusinessCopy

One of many snapshot methods included in Advanced Client. BusinessCopy is for making snapshots of client data on HP disk arrays.

Cache

Copy-on-write snapshot methods need a separate working area on disk during the lifetime of the snapshot, called a *cache*. The snapshot method uses the cache to store a copy of the client’s data blocks that are about to be changed by file system activity (see “[How Copy-on-Write Works](#)” on page 232 for a complete description). This cache must be a raw disk partition that does not contain valuable information: when using the cache, the snapshot method overwrites any data currently stored there.

Copy manager (see Third-Party Copy Device)

Copy-on-Write

In NetBackup Advanced Client, one of two types of supported snapshots (see also *mirror*). Unlike a mirror, a copy-on-write does not create a separate copy of the client’s data. It creates a block-by-block “account” that describes which blocks in the client data have changed and which have not, from the instant the copy-on-write was activated. This account is used by the backup application to create the backup copy. Other terms and trade names sometimes used for copy-on-write snapshots are space-optimized snapshots, space-efficient snapshots, and checkpoints.



Data movement

A copy operation as performed by a third-party copy device or NetBackup media server.

Data mover

The host or entity that manages the backup on behalf of the NetBackup client. This is either the NetBackup media server, a third-party copy device, or a NAS filer.

Disk group

A configuration of disks to create a primary-mirror association, using commands unique to the disks' vendor. See *mirror* and *volume group*.

Extent

A contiguous set of disk blocks allocated for a file and represented by three values: device identifier, starting block address (offset in the device) and length (number of contiguous blocks). The *mapping methods* in Advanced Client determine the list of extents and send the list to the backup agent.

FastResync (VxVM)

Formerly known as Fast Mirror Resynchronization or FMR, VxVM FastResync performs quick and efficient resynchronization of mirrors. NetBackup's Instant Recovery feature utilizes FastResync to create and maintain a point-in-time copy of a production volume.

Fibre channel

A type of high-speed network composed of either optical or copper cable and employing the Fibre Channel protocol. NetBackup Advanced Client supports both arbitrated loop and switched fabric (switched fibre channel) environments.

File system

Has two different meanings.

- ◆ When referring to a product, such as the UFS (Sun Solaris) or VxFS (VERITAS) *file system*, it refers to the management and allocation schemes on which the entire file tree is structured.
- ◆ When referring to a particular component in a file tree, *file system* means a directory (with any subdirectories and files) that is attached to the UNIX file tree by means of the `mount` command. When a file system is selected as an entry in the NetBackup Backup Selections list, this definition applies.

Instant Recovery

A restore feature of a disk snapshot of a client file system or volume. Client data can be rapidly restored from the snapshot, even after a system reboot.

Mapping

The process of converting a file or raw device (in the file system or Volume Manager) to absolute physical disk addresses or extents for use by backup agents on the network. NetBackup Advanced Client uses the *VxMS* library to perform file mapping.

Mapping methods

A set of routines for converting logical file addresses to absolute physical disk addresses or extents. NetBackup Advanced Client includes support for file-mapping and volume-mapping methods.

Mirror

- ◆ A disk that maintains an exact copy or duplicate of another disk. A mirror disk is often called a *secondary*, and the disk that it copies is called the *primary*. All writes to the primary disk are also made to the *mirror* (or secondary) disk.
- ◆ A type of snapshot captured on a mirror disk. At an appropriate moment, all further writes to the primary disk are held back from the mirror, thus causing the mirror to be “split” from the primary. As a result of the split, the mirror becomes a snapshot of the primary. The snapshot can then be backed up.

NetBackup Media Server method

An offhost backup method in which data movement is performed by a NetBackup media server.

Offhost backup

The off-loading of backup processing to a separate backup agent executing on another host. NetBackup Advanced Client provides the following offhost backup options: Alternate Client, NetBackup Media Server, Third-Party Copy Device, and Network Attached Storage.

Primary disk

In a primary-mirror array configuration, the primary is the disk on which client data is stored and is directly accessed by client applications. An exact duplicate of the primary disk is the mirror.

Raw partition

A single section of a raw physical disk device occupying a range of disk sectors, without a file system or other hierarchical organization scheme (thus, a “raw” stream of disk sectors). On some operating systems, such as Solaris and HP-UX, this is different from a block device over which the file system is mounted.

Recovery Manager (RMAN)

Oracle's backup and recovery program. RMAN performs backup and restore by making requests to a NetBackup shared library.



RMAN Proxy Copy

An extension to the Oracle®i Media Management API which enables media management software such as NetBackup to perform data transfer directly.

SAN (Storage Area Network)

A Fibre Channel-based network connecting servers and storage devices. The storage devices are not attached to servers but to the network itself, and are visible to all servers on the network.

Secondary disk

See *mirror*.

ShadowImage

One of many snapshot methods included in Advanced Client. ShadowImage is for making snapshots of client data on Hitachi disk arrays.

Snapshot

A point-in-time disk version of the data prior to backup. A *snapshot* is created very rapidly, causing minimal impact on other applications. There are two basic types: *copy-on-write* and *mirror*.

Snapshot method

A set of routines for creating a snapshot. You can select the method, or let NetBackup select it when the backup is started (auto method).

Snapshot mirror

A disk mirror created by the VERITAS Volume Manager (VxVM). This is an exact copy of a primary volume at a particular moment, reproduced on a physically separate device.

Snapshot source

The entity (file system, raw partition, or logical volume) to which a snapshot method is applied. NetBackup automatically selects the snapshot source based on the entries in the policy's Backup Selections list.

Snapshot Volume

A mirror that has been split from the primary volume or device, and made available to users. Snapshot volumes are created by the VERITAS Volume Manager (VxVM) as a point-in-time copy of the primary volume. Subsequent changes made to the primary volume are recorded in the Data Change Log and can be used later to resynchronize with the primary volume by means of VxVM FastResync. Only the changes made while the snapshot volume was detached from the primary would be applied to the snapshot volume to make it identical to the primary volume again.

Standard device

Refers to the primary disk in an EMC primary-mirror disk array (see *primary disk*).

Storage Checkpoint (VxFS)

Provides a consistent and stable view of a file system image and keeps track of modified data blocks since the last checkpoint. Unlike a mirror, a VxFS Storage Checkpoint does not create a separate copy of the primary or original data. It creates a block-by-block account that describes which blocks in the original data have changed and which have not, from the instant the checkpoint was activated.

A Storage Checkpoint stores its information in available space on the primary file system, not on a separate or designated device. (Also, the `ls` command does not list Storage Checkpoint disk usage; you must use the `fsckptadm list` command instead.)

Third-Party Copy Device

This term has two meanings:

- ◆ A backup agent on the SAN that operates on behalf of backup applications. The *third-party copy device* receives backup data from a disk attached to Fibre Channel and sends it to a storage device, using the SCSI Extended Copy command. The third-party copy device is sometimes called a copy manager, third-party copy engine, or data mover. In SAN hardware configurations, a third-party copy device can be implemented as part of a bridge, router, or storage device. The third-party copy device may or may not be the device to which the storage units are connected.
- ◆ An offhost backup method in NetBackup Advanced Client that allows backups to be made by means of a backup agent on the SAN.

TimeFinder

One of many snapshot methods included in Advanced Client. TimeFinder is for making snapshots of client data on EMC disk arrays.

UFS file system

The UNIX File System (UFS), which is the default file system type on Sun Solaris. The UFS file system was formerly the Berkeley Fast File System.

VxMS (VERITAS Federated Mapping Services)

A library of routines (methods) used by NetBackup Advanced Client to obtain the physical addresses of logical disk objects such as files and volumes.

Volume

A virtual device configured over raw physical disk devices (not to be confused with a NetBackup *Media Manager* volume). Consists of a block and character device.



If a snapshot source exists over a volume, NetBackup automatically uses a volume *mapping* method to map the volume to physical device addresses. Any of the Advanced Client *snapshot methods* can be used when backing up client data configured over volumes.

Volume group

A logical grouping of disks, created with the VERITAS Volume Manager, to allow more efficient use of disk space.

VxFS

The VERITAS extent-based File System (VxFS), designed for high performance and large volumes of data.

VxVM

The VERITAS Volume Manager (VxVM), which provides logical volume management that can also be used in SAN environments.

Installation

This chapter explains how to install NetBackup Advanced Client software on UNIX and Windows platforms.

Prerequisites

- ◆ NetBackup Enterprise server 6.0 or later must be installed on the master and media servers. For performing local backups, the master/media server can be running any supported UNIX or Windows platform. For a list of platforms supported by NetBackup Advanced Client, including those supported for offhost data mover backups, refer to “[Advanced Client Assistance](#)” on page xvi.
- ◆ NetBackup 6.0 or later client software must be installed on clients. For AIX and Linux clients, the client data must be in a VxFS file system.

Note On the NetBackup client, the base NetBackup software and the Advanced Client software must be at the same level.

- ◆ For Instant Recovery using the VxFS_Checkpoint method, the VxFS File System with the Storage Checkpoints feature must be installed on clients.



Installing Advanced Client On UNIX

Loading From Media

Note If you are installing in a cluster environment, you must freeze the active node before you begin the installation process so that migrations do not occur during installation. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

1. Log in as root on the NetBackup server.
2. In a separate window on the master server, make sure a valid license key for NetBackup Advanced Client has been installed. To do this, enter the following command to list and add keys:

```
/usr/opensv/netbackup/bin/admincmd/get_license_key
```

3. Insert the options CD-ROM containing the Advanced Client software in the drive.
4. Change your working directory to the CD-ROM directory:

```
cd /cd_rom_directory
```

Where *cd_rom_directory* is the path to the directory where you can access the CD-ROM. On some platforms, it may be necessary to mount this directory.

5. To install Advanced Client software on the NetBackup server, execute the following:

```
./install
```

Since other NetBackup products are included on the CD-ROM, a menu appears.

6. Select **NetBackup Add-On Product Software**.

A second menu appears. Select **NetBackup Advanced Client**.

7. Enter **q** to quit selecting options. When asked if the list is correct, answer **y**.

NetBackup Advanced Client software is installed in

```
/usr/opensv/netbackup/vfms/hardware/os/version/
```

Where:

- ◆ *hardware* is Solaris, HP9000-700, HP9000-800, RS6000, Linux
- ◆ *os* is Solaris8, Solaris9, Solaris10, HP-UX11.00, HP-UX11.11, AIX5, RedHat2.4



- ◆ *version* is a six digit number representing the NetBackup version

8. In a clustered environment, the above steps must be done on each node in the cluster.

Note If you are installing in a cluster environment, unfreeze the active node after the installation completes. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running

Note The NetBackup 6.0 client software must be installed on the clients before performing the next procedure. For instructions, refer to the *NetBackup Installation Guide for UNIX*. Note that the base NetBackup client software and the Advanced Client client software *must be at the same level*.

Distributing Advanced Client Software to UNIX Clients

Note If the server is a Solaris, HP, AIX, or Linux system, and the server is also a client, you do not have to distribute Advanced Client software to the local server. This distribution is done automatically when you load the software from the media. However, you must distribute the software as explained below to all other UNIX clients that will be using Advanced Client.

Note Advanced Client software cannot be distributed to clients by means of the NetBackup Administration Console.

Note If installing in a clustered environment, you must do the following on the active node of the cluster.

You should also perform this procedure if you are installing a software upgrade (patch).

Execute the following as the root user on the NetBackup server.

1. Check the Activity Monitor (or use the `bpdjobs` command) to make sure there are no backup or restore jobs running on the clients to which you want to distribute Advanced Client software. If a backup or restore job is listed as active or queued, wait until it is done.
2. If `bprd` is running on the master server (use `/usr/openv/netbackup/bin/bpps` to find out), terminate `bprd` by executing the following:

```
/usr/openv/netbackup/bin/admincmd/bprdreq -terminate
```



3. You can distribute the Advanced Client software to Solaris, HP, AIX, and Linux clients in either of two ways:

Note If distributing the Advanced Client software to clients located in a cluster, specify the host names of the individual nodes (not virtual names) in the list of clients.

- a. Distribute the software to all currently defined clients by executing the following command:

```
/usr/opensv/netbackup/bin/update_clients -Install_ADC
```

In each backup policy, clients must be configured with their correct OS type. For example, an HP client running HP-UX11.11 must be configured as such.

Note If you are installing an Advanced Client patch and your clients are not included in an existing NetBackup policy, Advanced Client software will not be distributed to the clients or alternate clients. You must create a file containing the client list and use the `update_clients` command as explained in the next step.

- b. Distribute the software to specific clients.

- ◆ Create a file that lists the specific clients. For each client, enter a line in this file containing the following three parameters:

```
hardware_type operating_system clientname
```

For example:

```
Solaris Solaris8 othersparc
```

or

```
HP9000-800 HP-UX11.11 myhp
```

or

```
Linux RedHat2.4 myredhat
```

or

```
RS6000 AIX5 myaix
```

- ◆ Execute the following command (all on one line):

```
/usr/opensv/netbackup/bin/update_clients -Install_ADC -ClientList file
```

Where *file* is the name of the file that you created in the previous step.

4. Start the NetBackup daemon as the root user on the master server by executing:

```
/usr/opensv/netbackup/bin/initbprd
```



Installing Advanced Client On Windows


For Windows, NetBackup Advanced Client software is automatically installed with the core NetBackup server and client product. For the NetBackup installation procedure, see the *NetBackup Installation Guide for Windows*.

You must add a valid license key for Advanced Client, as follows.

On each Windows master server:

Note If you are installing in a cluster environment, you must freeze the active node before you begin the installation process so that migrations do not occur during installation. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

1. Log in.
2. Add a license key for Advanced Client:
 - a. From the NetBackup Administration window, choose **Help**.
 - b. From the **Help** menu, select **License Keys ...**.

The NetBackup License Keys window appears. Existing keys are listed in the lower part of the window.
 - c. To register a new key, click the star icon  to open the Add a New License Key dialog. Type the new license key in the **New license key** field and click **Add**.

The new license key appears in the lower part of the dialog box.
3. In a clustered environment, the above steps must be done on each node in the cluster.

Note If you are installing in a cluster environment, unfreeze the active node after the installation completes. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.



Distributing Client Software in Mixed-Platform Environments

NetBackup Advanced Client software can be “pushed” from server to client in an all-UNIX environment only: server and clients must be running UNIX. Please see [“Distributing Advanced Client Software to UNIX Clients”](#) on page 31.

For Windows PCs, Advanced Client software is automatically installed with the base NetBackup software, for server and clients. Please see the *NetBackup Installation Guide for Windows*, and [“Installing Advanced Client On Windows”](#) on page 33 of this Advanced Client guide.

For mixed environments, please note the following:

- ◆ If you have a Windows server with UNIX clients (no UNIX server), you must set up a NetBackup UNIX server to distribute the software to your clients.
- ◆ If you have a UNIX server with Windows clients, you must install the client software on the Windows client machines individually, from the NetBackup CD-ROM. Please see [“Installing Advanced Client On Windows”](#) on page 33.



Creating Log Directories

During backup and restore, Advanced Client messages are written to several log directories on the NetBackup server and client. You must create these directories manually. Refer to [“Gathering Information and Checking Logs”](#) on page 214 for help creating log directories.

Upgrading from Earlier Releases

You can upgrade to the 6.0 version of NetBackup only if you currently have a NetBackup 5.x or 5.x MP version of the software installed. If you have an earlier version (for example NetBackup 4.5), you cannot upgrade directly to NetBackup 6.0. You must first upgrade to NetBackup 5.x and then upgrade.

Uninstalling Advanced Client

Server Uninstall - UNIX

On the server where you initially loaded the NetBackup Advanced Client or other software, do the following.

Note The following procedure results in total removal of the Advanced Client software.

Note If you are uninstalling in a cluster environment, you must first freeze the active node so that migrations do not occur during the uninstall. For information about freezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

On Solaris:

1. Check the Activity Monitor in the NetBackup Administration Console to make sure no NetBackup backups are active or running (the **State** field should read Done).
2. To remove the NetBackup Advanced Client package, execute the following:

```
pkgrm VRT$nbadc
```
3. See the following note on cluster environments.



On HP, AIX, Linux:

1. Check the Activity Monitor in the NetBackup Administration Console to make sure no NetBackup backups are active or running (the **State** field should read Done).
2. Execute the following:

```
/usr/opensv/netbackup/bin/install_adc -d  
rm -rf /usr/opensv/netbackup/vfms
```

Note If you are uninstalling in a cluster environment, unfreeze the active node after the uninstall. For information about unfreezing a service group, see the clustering section in the *NetBackup High Availability System Administrator's Guide* for the cluster software you are running.

Client Uninstall - UNIX

1. On the master server, check the Activity Monitor in the NetBackup Administration Console to make sure no NetBackup Advanced Client backups are active or running for the client (the **State** field should read Done).
2. Execute the following command to deinstall the NetBackup Advanced Client software on the client:

```
/usr/opensv/netbackup/bin/install_adc -d
```

Server/Client Uninstall - Windows

1. On the master server, check the Activity Monitor in the NetBackup Administration Console to make sure no NetBackup Advanced Client backups are active or running for the client (the **Job State** field should read Done).
2. Perform the basic uninstall procedure described in the *NetBackup Installation Guide for Windows*. Note: this uninstalls all of NetBackup.

SAN Configuration for Advanced Client

Due to the complex and rapidly changing nature of SAN configuration, this chapter describes SAN issues pertaining to NetBackup Advanced Client only. Please note the following assumptions:

- ◆ You have considerable technical expertise in both SAN and NetBackup configuration.
- ◆ Your hardware environment is already configured and functional, including switches, hubs, optional bridges or third-party copy devices, robots, tape drives, and disk arrays.
- ◆ Tape devices are visible to the NetBackup media server, and all passthru paths exist for tape devices and third-party copy devices.

Note This chapter applies to the NetBackup Media Server and Third-Party Copy Device backup methods only. If your backup policies are not using either of these methods, you may skip this chapter.

This chapter includes the following topics:

- ◆ [SAN Configuration Diagram](#)
- ◆ [Media Server/Third-Party Copy Requirements](#)
- ◆ [Configuration Flowcharts](#)
- ◆ [Verify NetBackup Access to SAN Devices](#)
- ◆ [Solaris only: Configure HBA Drivers](#)
- ◆ [Create Backup Configuration Files](#)



SAN Configuration Diagram

The following diagram shows the devices and configuration files described by the procedures in this chapter. This diagram shows devices configured behind a third-party copy device (bridge) as well as directly attached to the SAN.


Backup configuration files on **media server**:

3pc.conf file:

Contains tape and client disk info used by third-party copy device.

mover.conf file:

Identifies the third-party copy device.



Supported Peripherals

A complete list of Advanced Client supported peripherals can be found on the VERITAS support web site. For instructions, refer to “[Advanced Client Assistance](#)” on page xvi.

Media Server/Third-Party Copy Requirements

The configuration requirements for NetBackup Media Server and Third-Party Copy Device backup methods are as follows:

- ◆ The information needed for the NetBackup Media Server method is obtained by means of the `bptpcinfo` command described later in this chapter.
- ◆ The information needed for the Third-Party Copy Device method requires the `bptpcinfo` and `bpmoverinfo` commands. Depending on your devices, you may need the following:
 - ◆ The VERITAS CommandCentral Storage product (or SANPoint Control) for locating the world-wide port name and luns. If you have CommandCentral Storage or SANPoint Control, you can use the `bpSALinfo` command to look up world-wide name and lun information. See the note below.
 - ◆ The instructions provided with your HBA and bridge/router/third-party copy device. The VERITAS support web site also contains information to help you configure devices (see “[Advanced Client Assistance](#)” on page xvi for instructions).

Note To use the NetBackup `bpSALinfo` command with CommandCentral Storage, you need the following:

- A host running CommandCentral Storage Server, licensed for Operations Module clients.
- On each NetBackup client, a CommandCentral Storage Operations Module Agent.

See “[Create the 3pc.conf File](#)” on page 69 for instructions on using the `bpSALinfo` command.




Diagram for NetBackup Media Server

In this backup method, the NetBackup media server handles the backup processing and sends the backup data over Fibre Channel to the storage device.

Note If you have a multi-ported SCSI disk array, a fibre channel SAN is not required. See “[Offhost Backup Without a SAN \(UNIX Only\)](#)” on page 17.

NetBackup Media Server Backup Method

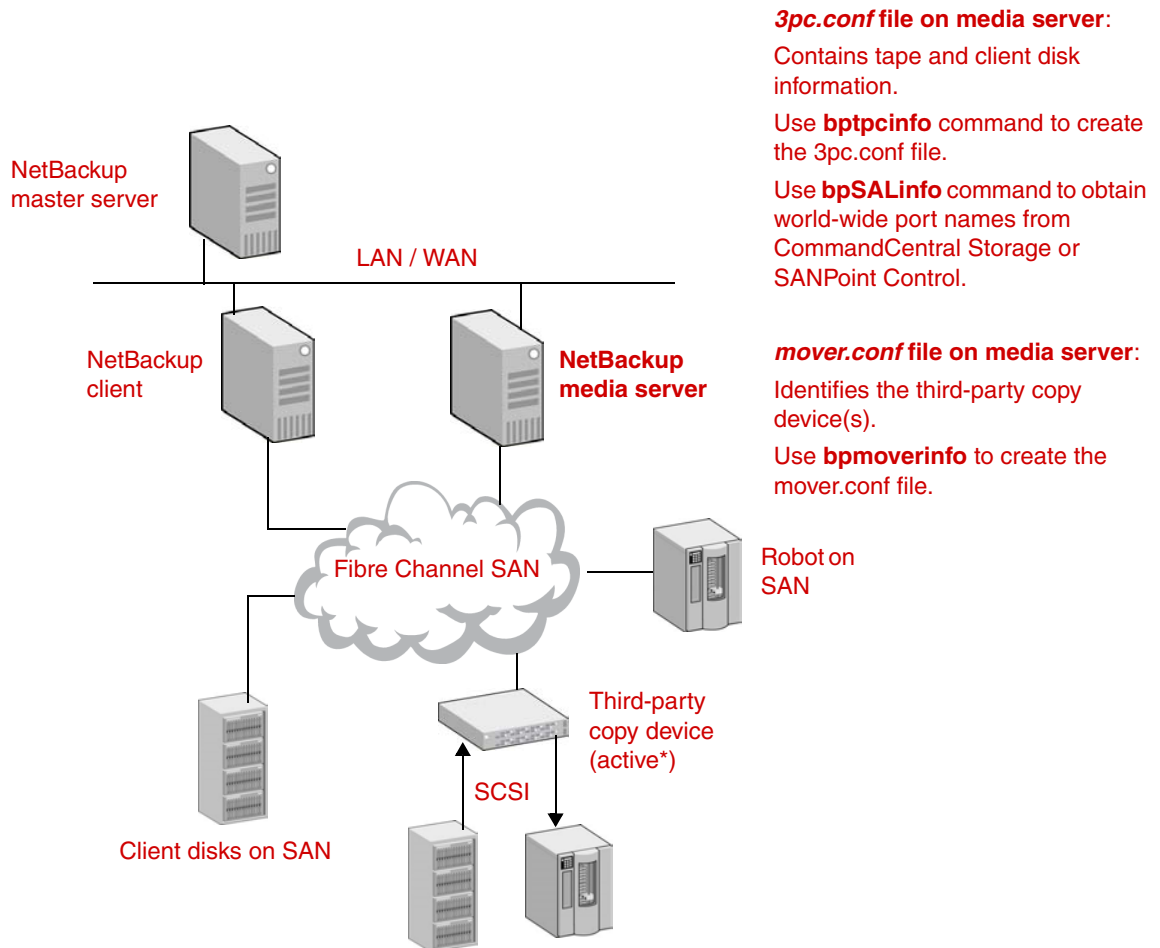


**passive* means the third-party copy device allows media server access to the disks/tapes but does not perform SCSI Extended Copy commands.

Diagram for Third-Party Copy Device

In this backup method, a third-party copy device performs the I/O processing of the backup (data movement). The `3pc.conf` file describes the disks to be backed up and the tape devices for data storage. Be sure to use the flowcharts under “[Configuration Flowcharts](#)” on page 43.

Third-Party Copy Device Backup Method



*active means the third-party copy device performs SCSI Extended Copy commands to move the data.

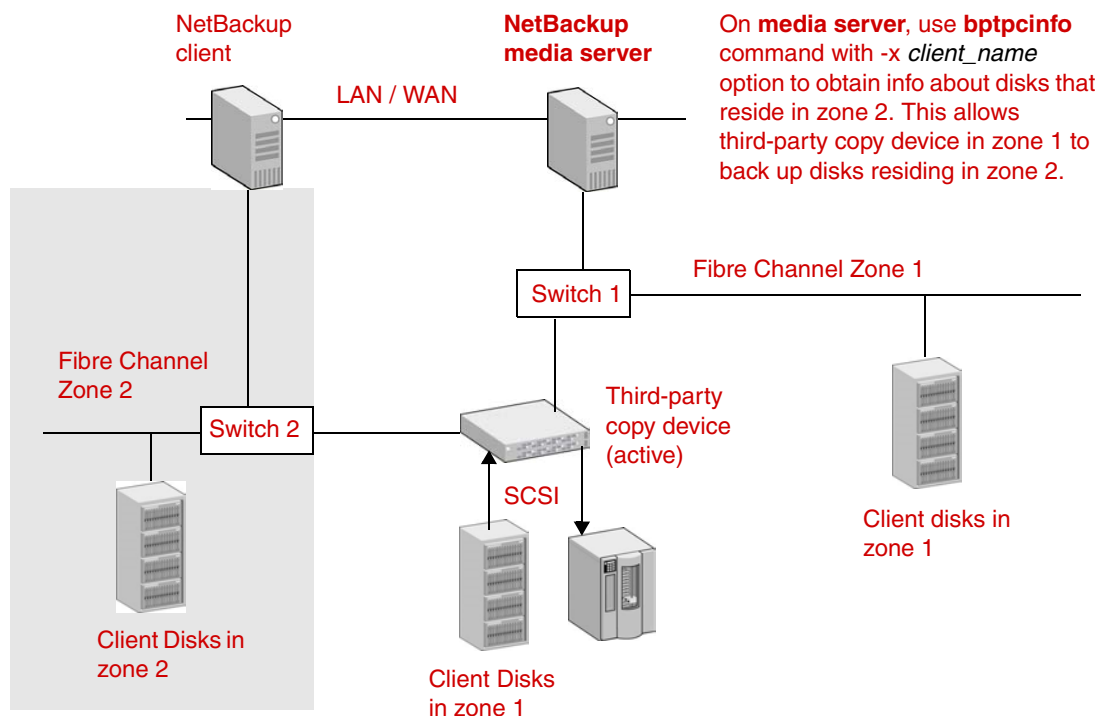


Diagram for Third-Party Copy Device - Remote

In this configuration, the NetBackup media server and the disks containing the client data are on different fibre-channel networks (due to zoning or LUN-masking). The media server can communicate with the NetBackup client by means of the LAN, but does not have access to the client's disks located on a different fibre channel network (or zone). In this case, the 3pc.conf file must be modified with the `bptpcinfo` command using the `-x client_name` option, to include information about the client disks.

Note If all devices support identification descriptors (E4 target), NetBackup automatically creates a complete 3pc.conf file and you do not need to run the `bptpcinfo` command.

Third-Party Copy: Remote



Configuration Flowcharts

The following four charts show the process for setting up configuration files for Media Server or Third-Party Copy backup. Instructions are included later in this chapter.

Chart I: Verify SAN Device Visibility for NetBackup Media Server

For instructions, see “Verify NetBackup Access to SAN Devices” on page 47.

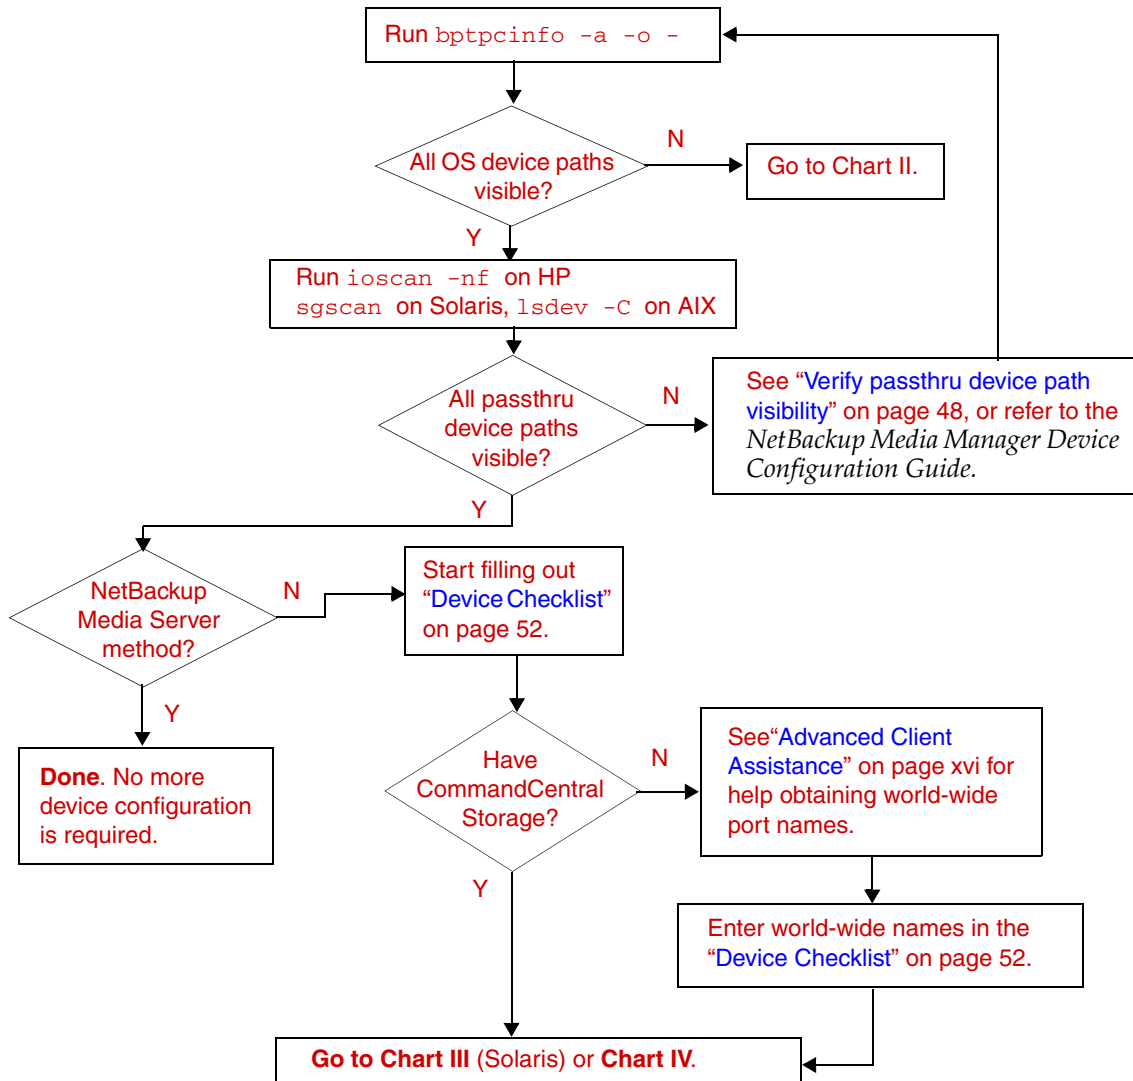


Chart II: Verify OS Device Paths Visibility

For instructions, see “[Making OS device paths visible](#)” on page 47.

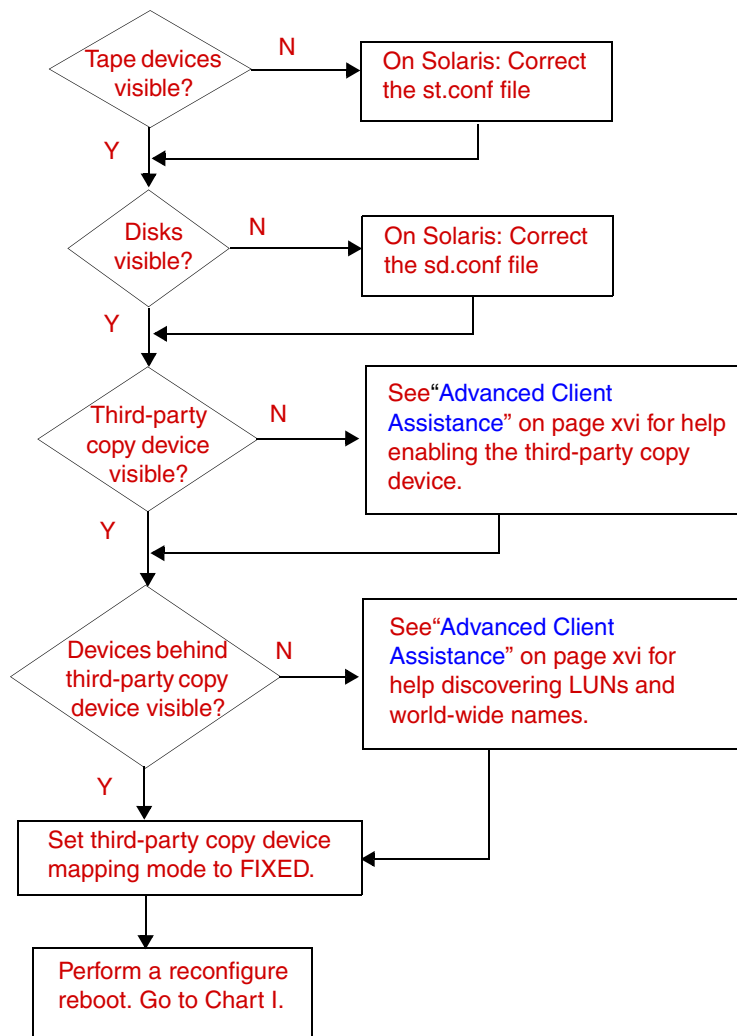


Chart III: Solaris only: Configure HBA Drivers

For instructions, see “[Solaris only: Configure HBA Drivers](#)” on page 53.

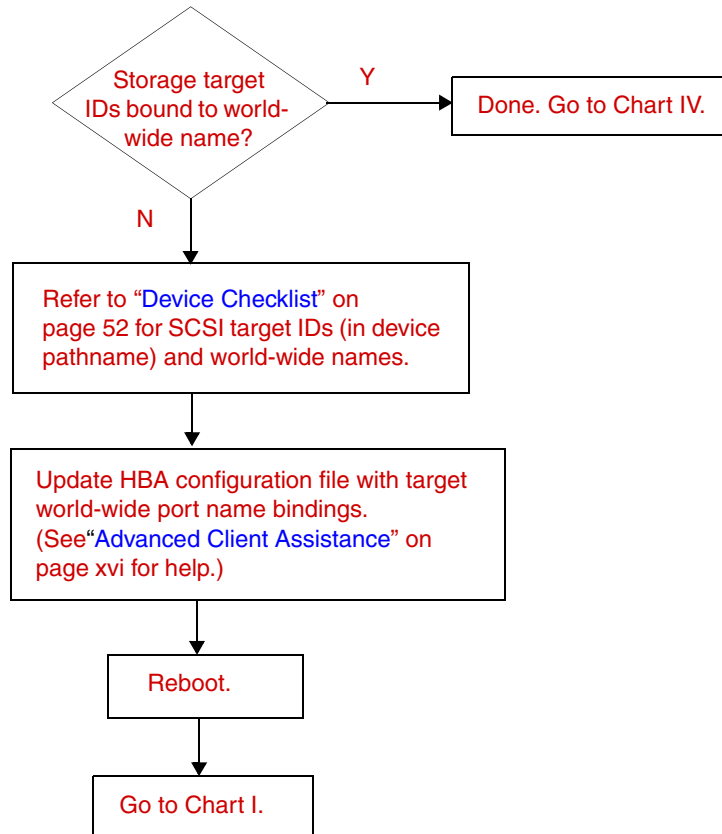
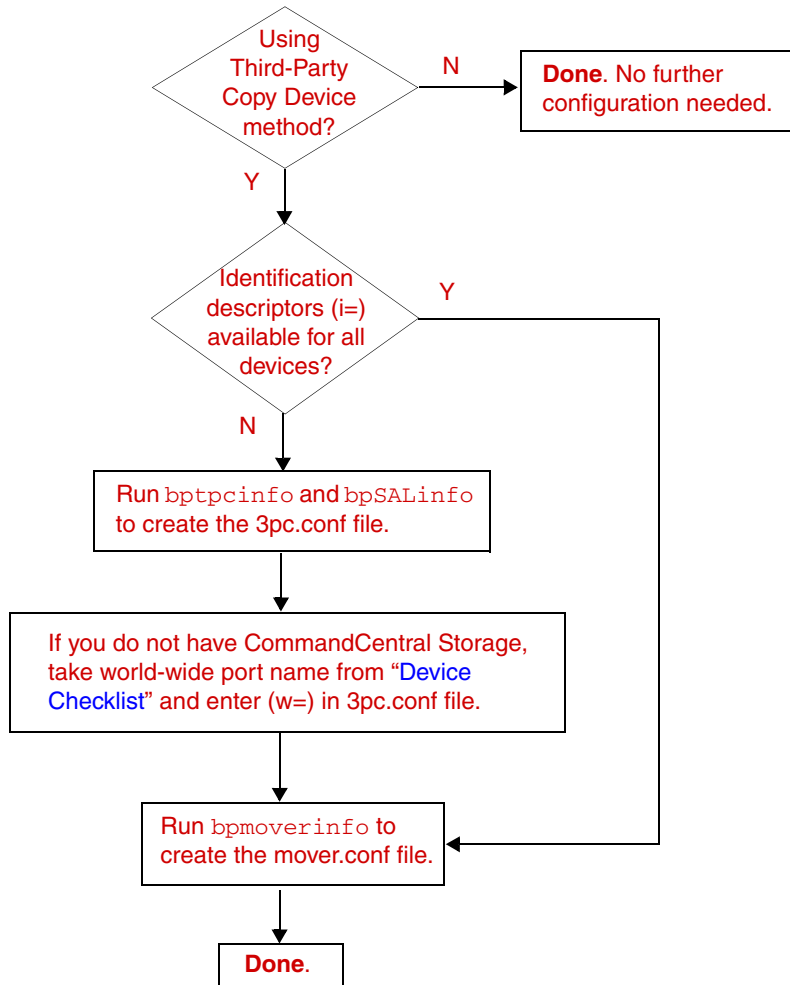


Chart IV: Create the Backup Configuration Files

For instructions, see “Create Backup Configuration Files” on page 54.



Verify NetBackup Access to SAN Devices

Note It is assumed that NetBackup and all device drivers are installed, and that devices are properly connected and powered up.

▼ Verify OS device path visibility

1. On the media server, run the `bptpcinfo` command.

The following sends the output to the screen, using `-o -` (note the space before the final hyphen).

```
/usr/opensv/netbackup/bin/bptpcinfo -a -o -
```

The following sends the output to a file:

```
/usr/opensv/netbackup/bin/bptpcinfo -a -o output_file_name
```

2. Examine the `bptpcinfo` output to see if your OS device paths are listed. If all devices are listed, go to [step 8](#) for HP, [step 9](#) for AIX, or to [step 10](#) for Solaris.

▼ Making OS device paths visible

3. **For Solaris:** If your tape devices are not listed in the `bptpcinfo` output, make sure you have target and LUN values for each tape device in the `st.conf` file.
4. **For Solaris:** If your disks are not listed in the `bptpcinfo` output, make sure you have target and LUN values for each disk in the `sd.conf` file.
5. If the devices behind the bridge (or third-party copy device) are not listed in the `bptpcinfo` output, or if the third-party copy device is not enabled for third-party copy data movement, see the VERITAS support website for assistance (see [“Advanced Client Assistance”](#) on page xvi).
6. On the bridge or third-party copy device, set the address mapping mode to FIXED. This prevents the addresses from changing when the devices are reset. For help configuring third-party copy devices, see the VERITAS support website (see [“Advanced Client Assistance”](#) on page xvi).
7. Enter the following to reboot the operating system on the media server:

Solaris:

```
reboot -- -r
```

HP and AIX:

```
reboot
```



▼ Verify passthru device path visibility

- 8. For HP:** Enter the following to list all passthru devices:

```
ioscan -nf
```

- a.** If all devices now appear, enter the following to regenerate HP special files:

```
insf -e
```

Then go to [step 11](#) on page 49.

- b.** If some devices do not appear in the `ioscan` output, check hardware connections to the devices that are not appearing. Then repeat [step 8](#).

Note On HP 11.00, there is a limit of eight devices per target. For instance, if you have a JBOD disk array consisting of ten disks, and the array is connected to a bridge, it may be that only the first eight disks in the array are accessible.

- 9. For AIX:**

- a.** Enter the following to list all passthru devices and create the paths:

```
cfgmgr
```

- b.** Enter the following to list the results:

```
lsdev -C
```

If all devices now appear, go to [step 11](#) on page 49.

- c.** If some devices do not appear in the output, check hardware connections to the devices that are not appearing. Then repeat [step 9](#).

- 10. For Solaris:**

- a.** Perform an `sgscan` to list all passthru devices. Check for proper output and recognition of devices.

Here is sample output from `sgscan`:

```
/dev/sg/c0t6l4: Tape (/dev/rmt/2): "QUANTUM DLT7000"  
/dev/sg/c0t6l5: Changer: "HP C6280-7000"
```

- b.** If tape devices still do not show up, make sure you have entries for all SCSI target and LUN combinations in the `sg.links` and `sg.conf` files. Refer to the *Media Manager Device Configuration Guide*, Chapter 2, under “Understanding the SCSI Passthru Drivers.”

- ◆ If tape devices are fibre attached, make sure you have entries for the tape devices in the above files.
- ◆ If tape devices are behind a bridge (or third-party copy device), make sure you have entries for the tape devices AND for the bridge/third-party copy device.

For an example, refer to [“Solaris only: Example for sg.links, sg.conf, and st.conf files”](#) on page 51.

If you are unsure how to acquire the SCSI target and LUN values for your configuration, see [“Advanced Client Assistance”](#) on page xvi for help with particular devices. For instance, if your tape drives are configured behind a bridge, router or other fibre-channel device, you may need to telnet into the device to determine the target ID and LUN for each tape drive.

- c. When finished updating the `sg.links`, `sg.conf`, and `st.conf` files, remove the old sg configuration:

```
rm /kernel/drv/sg.conf
rem_drv sg
```

- d. Run the `/usr/opensv/volmgr/bin/driver/sg.install` script to copy the files into the correct locations.
- e. Copy the `sg.links` and `sg.conf` files (in `/usr/opensv/volmgr/bin/driver`) to another location, for future reference. Whenever NetBackup is re-installed, these files in `/usr/opensv/volmgr/bin/driver` are overwritten.

11. Run the `bptpcinfo` command again to see which devices are now visible to the media server. Repeat at [step 2](#) if any of your SAN devices are not showing up in the `bptpcinfo` command output.
12. If the offhost backup method is NetBackup Media Server, no more device configuration is required. You can skip the rest of this chapter.
13. When all devices are listed in the `bptpcinfo` command output, use that information to fill in the device pathname (p=), serial number (s=), and LUN (l=) in the [“Device Checklist”](#) on page 52 for each device.

▼ Determine the world-wide name for each device

There are two options:

- ◆ Use VERITAS CommandCentral Storage, formerly SANPoint Control ([step 14](#)).
- ◆ Otherwise, refer to [“Advanced Client Assistance”](#) on page xvi for help determining the world-wide name of your devices. Then go to [step 15](#).



14. You can use VERITAS CommandCentral Storage (or SANPoint Control) to determine the world-wide port names (storage ports) for the devices.

Managing

Reporting

Monitoring

Topology

Task Status

Tools

Settings

Disabled

SummaryStorageApplicationsDatabasesHosts and HBASANSGroups

Connected To
kalashree.vxindia.veritas.cor

Alert Summary
14021

Global Alerts
Custom Reports
Hosts Summary
Primary Host
Enclosures Summary
HBA Replacement Tool
Zone Builder
Storage Provisioning Tool
Switches Summary
LUN Builder
LUN Query Tool
Snapshot Editor
Switches

Managing SummaryUnenclosed DevicesHewlett-Pa 32

Generic Device: Hewlett-Pa 32

Create AttributeGo

OverviewLUNsConnectivityZoningTopologyAttributes

Claimed Storage

Name	Host	HBA	Storage Port	LUN
/dev/rmt/3	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	Drive1
/dev/sg/c2t310	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	Drive1
/dev/sg/c2t311	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 1
/dev/sg/c2t312	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 2
/dev/sg/c2t313	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 3
/dev/sg/c2t314	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 4
/dev/sg/c2t315	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 5
/dev/sg/c2t316	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 6
/dev/sg/c2t317	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 7
/dev/sg/c2t318	susoe02.ilabeast.nsmq	HBA 20.00.00.e0.8b.06.b5.1b	Port 50.06.0b.00.00.1e.01.fc	LUN 8
Total: 17				

15. Update the “Device Checklist” on page 52 with the world-wide port names of your devices.

Note It is important to record this information. It will be needed again, to complete the configuration.

16. For Solaris: continue with “Solaris only: Configure HBA Drivers” on page 53. For HP and AIX, continue with “Create Backup Configuration Files” on page 54.

Solaris only: Example for `sg.links`, `sg.conf`, and `st.conf` files

The following is an example for [step 10](#) on page 48. For the devices in this example, fibre channel LUNs 0, 1, 4, and 5 are needed for target (Loop ID) 6. In this example, LUN 0 is the third-party copy device, LUN 1 is the tape controller, and LUNs 4 and 5 are the tape drives.

- ◆ Add entries in the `/usr/opensv/volmgr/bin/driver/sg.links` file so that the necessary `/dev/sg/*` nodes are created.

Note that the target and LUNs in the address part of the `sg.links` entries are hexadecimal, but are decimal in the `sg/c\N0tmln` part of the entries. Also, make sure there are tabs between the columns, not spaces.

```
type=ddi_pseudo;name=sg;addr=6,0;      sg/c\N0t6l0
type=ddi_pseudo;name=sg;addr=6,1;      sg/c\N0t6l1
type=ddi_pseudo;name=sg;addr=6,4;      sg/c\N0t6l4
type=ddi_pseudo;name=sg;addr=6,5;      sg/c\N0t6l5
```

- ◆ Add additional target and LUN entries to the `/usr/opensv/volmgr/bin/driver/sg.conf` file.

```
name="sg" class="scsi" target=6 lun=0;
name="sg" class="scsi" target=6 lun=1;
name="sg" class="scsi" target=6 lun=4;
name="sg" class="scsi" target=6 lun=5;
```

- ◆ In the `/kernel/drv/st.conf` file, do the following:
 - ◆ Add (or un-comment) the appropriate drive entries in the `tape-config-list` section.

```
tape-config-list =
    "DEC      TZ89",          "DEC DLT",          "DLT7k-data";
```

- ◆ Then add (un-comment) the matching `data-property-name` entry:

```
DLT7k-data = 1,0x38,0,0x39639,4,0x82,0x83,0x84,0x85,2;
```

- ◆ For each tape drive, add a name entry to the `st.conf` file.

Here is an example name entry:

```
name="st" class="scsi" target=6 lun=4;
name="st" class="scsi" target=6 lun=5;
```

Make sure you have entries for all target and bus combinations for each device.



Device Checklist

Use this checklist or one like it to record information about each of your SAN devices. Some of this information is provided by the `bptpcinfo` command (such as device pathname and serial number), and some has to be obtained by other means as explained in these procedures. It is vital that the information be recorded accurately.

Type of Device (disk or tape)	Device pathname used by UNIX host (p=)	Serial number (s=)	LUN (l=)	World-wide port name (w=)



Solaris only: Configure HBA Drivers

Fibre channel devices should be bound to specific SCSI target IDs by modifying the driver configuration files for your host bus adapter (HBA). This binding ensures that the host HBA and the third-party copy device are in agreement as to the target and LUN values for each device. The binding also ensures that the target ID does not change after a system reboot or after a fibre-channel reconfiguration. If the target ID changes, the backup configuration files (3pc.conf, mover.conf) will also be incorrect and will have to be recreated.

The binding process is unique to each vendor and product. For assistance, refer to the documentation provided for your HBA, or to the VERITAS support website. (See [“Advanced Client Assistance”](#) on page xvi.) The binding requires the fibre-channel world-wide port name.

Note Each time a device is added or removed, the binding must be updated to reflect the new configuration.

If storage device SCSI target IDs are bound to world-wide port names in your HBA configuration file, skip this section and go to [“Create Backup Configuration Files”](#) on page 54.

▼ To configure HBA drivers on the media server:

1. If storage device target IDs are not already bound to world-wide port names, refer to your [“Device Checklist”](#) on page 52 (filled out in the previous procedure) for the world-wide names. Use the world-wide names to make the binding for each device.
2. Update your HBA configuration file by binding all SCSI device target IDs to their associated world-wide port name.

For assistance with your particular HBA file, see [“Advanced Client Assistance”](#) on page xvi.

3. Reboot the media server (`reboot -- -r`).
4. To ensure device visibility, repeat the steps described under [“Verify NetBackup Access to SAN Devices”](#) on page 47.

When you are finished, the `btpcinfo` command should list device pathnames and serial numbers for all of your devices. Update the [“Device Checklist”](#) with those values if needed.

5. Continue with [“Create Backup Configuration Files.”](#)



Create Backup Configuration Files

This section pertains to the Third-Party Copy Device method only. If you are using any other backup method, you can skip the rest of this chapter.

For the Third-Party Copy Device method, you must create the following file on the media server:

```
/usr/openv/volmgr/database/mover.conf
```

If CommandCentral Storage (or SANPoint Control) does not support your environment and some devices (such as disks, tapes, or third-party copy devices) do NOT support identification descriptors (E4 target), you must add world-wide port name information for those devices in this file on the media server:

```
/usr/openv/volmgr/database/3pc.conf
```

Note At the start of the backup, NetBackup creates a 3pc.conf file if one does not exist. If all devices support identification descriptors, you do not need to create or edit the 3pc.conf file. You can skip to “[mover.conf Description](#)” on page 59 and to “[Create the mover.conf File](#)” on page 72.

The 3pc.conf and mover.conf Files: An Overview

The NetBackup media server needs certain information about the devices available on the SAN in order to coordinate the backup. This information is provided in two files:

- ◆ **3pc.conf:**

Identifies the client disks on the SAN that can be backed up, and the robotic libraries/tape drives on which NetBackup can store the data. The NetBackup media server uses this information to access client disks when performing the backup. It also uses this information to generate the SCSI Extended Copy commands required by third-party copy devices.

- ◆ **mover.conf:**

Identifies the third-party copy devices. These are devices that can execute the SCSI Extended Copy commands. A variety of devices can be designed to operate SCSI Extended Copy, such as routers, bridges, robotic libraries, and disk arrays. The mover.conf file is needed for the Third-Party Copy Device backup method only, not for the NetBackup Media Server method.

3pc.conf Description

In the `3pc.conf` file, each SAN device needs a one-line entry containing several kinds of values. The values required depend on several factors (explained below). Typically, these include (but are not limited to) the device ID, host-specific device path, and serial number. One or more of the following are also required: the identification descriptor, logical unit number (LUN) and world-wide port name. See “[Determining Requirements](#)” on page 58.

Some of this information will be automatically discovered and filled in by the `bptpcinfo` command, as described under “[What bptpcinfo Automatically Provides](#)” on page 58. The procedure for using the `bptpcinfo` command is under “[Create the 3pc.conf File](#)” on page 69.

Example 3pc.conf file

Below is an example `3pc.conf` file, followed by descriptions of each field. This file is located in `/usr/opensv/volmgr/database`.

Example 3pc.conf

```
# devid [a=wwpn:lun] [c=client] [p=devpath] [P=clientpath] [s=sn] [l=lun] [w=wwpn] [W=wwpn] [i=iddesc]
0 p=/dev/rdisk/c0t0d0s2 s=FUJITSU:MAB3091SSUN9.0G:01K52665 l=0
1 p=/dev/rdisk/c0t10d0s2 s=FUJITSU:MAG3091LSUN9.0G:00446161 l=0
2 p=/dev/rdisk/c4t0d0s2 s=HP:OPEN-3:30436000000 l=0 a=500060E80276E401:0
3 p=/dev/rdisk/c4t1d0s2 s=FUJITSU:MAN3367MSUN36G:01X37938 l=0 a=100000E00221C153:0
4 p=/dev/rdisk/c4t3d0s2 s=HITACHI:OPEN-3-CM:20461000000 l=0 i=10350060E800000000000004FED00000000 a=50060E80034FED00:0
5 p=/dev/rdisk/c4t14d0s2 s=HITACHI:OPEN-9:60159003900 l=0 w=500060e802eaff12
6 p=/dev/rdisk/c4t0d1s2 s=HP:OPEN-3:30436000100 l=1 a=500060E80276E401:1 a=1111222233334444:0
7 p=/dev/rdisk/c4t0d2s2 s=HP:OPEN-3:30436000200 l=2 a=500060E80276E401:2
8 p=/dev/rdisk/c4t0d3s2 s=HP:OPEN-3:30436000300 l=3 a=500060E80276E401:3
9 p=/dev/rdisk/c4t0d4s2 s=HP:OPEN-3-CM:30436005100 l=4 a=500060E80276E401:4
10 p=/dev/rdisk/c4t0d5s2 s=HP:OPEN-3:30436002600 l=5 a=500060E80276E401:5
11 p=/dev/rdisk/c4t0d6s2 s=HP:OPEN-3:30436002700 l=6 a=500060E80276E401:6
12 p=/dev/rdisk/c4t0d7s2 s=HP:OPEN-3:30436002800 l=7 a=500060E80276E401:7
13 p=/dev/rdisk/c4t0d8s2 s=HP:OPEN-3:30436002900 l=8 a=500060E80276E401:8
14 p=/dev/rdisk/c4t1d1s2 s=FUJITSU:MAN3367MSUN36G:01X37958 l=1 a=100000E00221C153:1
15 p=/dev/rdisk/c4t1d2s2 s=FUJITSU:MAN3367MSUN36G:01X38423 l=2 a=100000E00221C153:2
16 p=/dev/rdisk/c4t1d3s2 s=FUJITSU:MAN3367MSUN36G:01X38525 l=3 a=100000E00221C153:3
17 p=/dev/rdisk/c4t1d4s2 s=FUJITSU:MAN3367MSUN36G:01X37951 l=4 a=100000E00221C153:4
18 p=/dev/rdisk/c4t1d5s2 s=FUJITSU:MAN3367MSUN36G:01X39217 l=5 a=100000E00221C153:5
19 p=/dev/rdisk/c4t3d1s2 s=HITACHI:OPEN-3-SUN:20461000300 l=1 i=10350060E800000000000004FED00000003 a=50060E80034FED00:1
20 p=/dev/rdisk/c4t3d2s2 s=HITACHI:OPEN-3-SUN:20461000400 l=2 i=10350060E800000000000004FED00000004 a=50060E80034FED00:2
```

The `3pc.conf` file can contain the following types of entries (keyword, if any, is in parentheses):



(comment line, or “ignore-rest-of-line”)

A number sign (#) can be used to comment out whatever occurs to the right of the # sign to the end of the line. You can use this to introduce a comment, or to block out an existing device-entry or portion of a line. For instance, if you remove a device from your configuration temporarily, you could comment out the corresponding line for that device in the 3pc.conf file, then remove the # later when the device is added back to the configuration.

device ID (devid)

A unique NetBackup number for the device. In the 3pc.conf file, the device ID numbers need not be in sequential order, but must be unique.

address (a=wwpn:lun)

The world-wide port name and lun as provided by the bpSALinfo command (see [step 3](#) on page 70 for information on this command). For a device that has multiple FC ports, there can be multiple a= entries.

client name (c=client)

The name of the remote client provided by the bptpcinfo command with the -x option (see [step 2](#) on page 69 for more information on this command).

device path (p=devpath)

The path to the Fibre Channel device. This entry must be specific to the media server on which the 3pc.conf file resides.

client path (P=clientpath)

The path to the remote client provided by the bptpcinfo command with the -x option.

serial number (s=sn)

The serial number of the device, of the form:

Vendor ID:Product ID:device serial number

Note The disk devices must support SCSI serial-number inquiries or page code 83 inquiries. If a page code inquiry returns an identification descriptor (i=) for a disk, the serial number is not required.

lun (l=lun)

The device's logical unit number. The *LUN* allows NetBackup to identify devices that are attached by SCSI connection to the third-party copy device, bridge, or other SAN device, or that are directly attached to the fibre channel.

world-wide port name (w=wwpn)

Note w=wwpn is allowed for backwards compatibility with 4.5 clients. If you run the bpSALinfo or bptpcinfo command, the w=wwpn entry will be converted to a=wwpn:lun.

The device's fibre channel *world-wide port name*, which identifies the device on the SAN. This is a 16-digit identifier, consisting of an 8-digit manufacturer name, and an 8-digit device name (numeric).

The following is an example message showing a world-wide name for a device, written to the /var/adm/messages log on the server. Note there are two versions of the world-wide name: the *node wwn* and *port wwn*. **For Advanced Client, use the port wwn.**

Example of World-Wide Name in /var/adm/messages log

```
Dec 12 16:07:19 sundog unix:fca-pci0: Target 1 : Port 0000e8 (1000005013b10619:2000005013b10619) online
```

NOTE!

The format of this message may vary depending on the host bus adapter card that is used.

This is the "Node WWN"

This is the "Port WWN"

On some devices, the world-wide port name can be found on the back of the device or in the boot-time messages written to the /var/adm/messages log on the NetBackup media server.

alternate world-wide port name (W=wwpn)

Note W=wwpn is allowed for backwards compatibility with 4.5 clients. If you run the bpSALinfo or bptpcinfo command, the W=wwpn entry will be converted to a=wwpn:lun.

A disk in a disk array can be assigned to multiple Fibre Channel ports. This is done, for instance, for load balancing or redundancy, to allow other devices on the SAN to access the same disk through different ports. The two ports allow NetBackup to select the port by which the storage device will access the disk.



In such a configuration, while `w=` specifies the first world-wide port name, `W=` specifies the second world-wide port name for the disk. (Note the uppercase “W” in `W=wwpn`.)

identification descriptor (`i=iddesc`)

When available, this value (up to 43 characters) identifies the device on the SAN. If available, the identification descriptor is automatically included in the `3pc.conf` file when you run the `bptpcinfo` command. See “[Determining Requirements](#)” for more information on this value.

Determining Requirements

The following determines which values are required in the `3pc.conf` file.

identification descriptor

The identification descriptor is optional, and is not supported by all vendors. To produce this descriptor, the device must support a page code inquiry with a type 2 or 3 descriptor of less than 20 bytes. The NetBackup `bptpcinfo` command (explained below) will detect the device’s identification descriptor and place it in the `3pc.conf` file if the identification descriptor is available.

Even when this descriptor is available, some third-party copy devices do not support its use.

Note If an identification descriptor is available and the third-party copy device supports it, the descriptor is used to identify the device on the SAN; in this case, there is no need for the LUN or world-wide name. To determine whether your third-party copy device supports identification descriptors, see “[Advanced Client Assistance](#)” on page xvi.

world-wide port name

If an identification descriptor is not available or the third-party copy device does not support identification descriptors, the device’s world-wide port name must be included in the `3pc.conf` file.

What `bptpcinfo` Automatically Provides

The NetBackup `bptpcinfo` command detects some or all of the device information needed for the backup and places that information in the `3pc.conf` file, as follows:

- ◆ The `bptpcinfo` command provides the device path, serial number, identification descriptor (if available), and the LUN.

- ◆ The `bptpcinfo` command does not provide the world-wide name. If you have CommandCentral Storage (or SANPoint Control), use `bpSALinfo` to obtain world-wide names.

What the Backup Methods Require

For the NetBackup Media Server backup method, the `bptpcinfo` command provides all the information you need (no manual editing is required).

The Third-Party Copy Device method requires more information for each disk. In some instances, the `bptpcinfo` command cannot gather all the information required.

mover.conf Description

The `/usr/openv/volmgr/database/mover.conf` file identifies the third-party copy devices that NetBackup can use for the Third-Party Copy Device backup method. This file is needed for the Third-Party Copy Device backup method only.

You can use the `bpmoverinfo` command to create the `mover.conf` file (see “[Create the mover.conf File](#)” on page 72). In most cases, the `bpmoverinfo` command makes the appropriate entry in the `mover.conf` file and no further configuration is needed. The next few sections describe the types of entries that can be made in `mover.conf`, in case manual configuration is needed.

Types of entries in mover.conf

Depending on your device configuration, the `mover.conf` file can consist of the following:

- ◆ the `passthru` driver device path
- ◆ the `/dev/rmt/device_name` of the third-party-copy-capable tape drive(s), or the `TAPE` keyword
- ◆ the `/dev/rdisk/device_name` of the third-party-copy-capable disk drive(s), or the `DISK` keyword


Please note the following:

passthru driver device path

You must enter a `passthru` path (`/dev/sg` for Solaris, `/dev/sctl` for HP-UX, or `/dev/ovpassn` for AIX) in the `mover.conf` file when the storage device for the backup is not “behind” (SCSI-connected to) the third-party copy device that will handle the backup.



When passthru path is required in mover.conf file



In this example, to use the third-party copy device to send the backup to **Tape 1**, the mover.conf file must include the passthru device path of the third-party copy device: `/dev/sg/c0t16l0`.


/dev/rmt/device_name of the tape

/dev/rdisk/device_name of the disk

You can enter the `/dev/rmt/device_name` of a tape or the `/dev/rdisk/device_name` of a disk in the mover.conf file when the tape or disk is behind a third-party copy device, or when the tape or disk has built-in third-party copy functionality. The third-party copy device in front of (or inside) the tape or disk drive is used for the backup.

Note The `/dev/rmt/device_name` path will be used if it matches the drive path that NetBackup selects for the backup. As a rule, this is not a problem, since the `bpmoverinfo` command detects all available third-party copy devices (and any tape devices behind them) and enters them in the mover.conf file. See [“Create the mover.conf File”](#) on page 72.

`/dev/rmt/device_name` in `mover.conf` file



In this example, to use a third-party copy device to send the backup to **Tape 2** or to **Tape 3**, the `mover.conf` file can specify the `device_name` of the tape drive: `/dev/rmt/2cbn` or `/dev/rmt/3cbn`. To use **Tape 1**, the `mover.conf` file would need the `passthru` device path of a third-party copy device.

Note To use a tape unit as a third-party copy device (such as Tape 3), a SCSI passthru device path must have been configured for that tape device.

TAPE and DISK keywords

For convenience, you can specify the `TAPE` keyword or `DISK` keyword in `mover.conf` instead of listing individual tape or disk `device_name` paths. In the above example, the `TAPE` keyword could be used to send the backup to either Tape 2 or Tape 3.

See “[Keywords in Mover File](#)” on page 64 for more information on these keywords.

For sites that have one third-party copy device

- ◆ The `mover.conf` file can consist of one line specifying the device by means of its device path:

For example, on Solaris:

```
/dev/sg/c6t110
```

On HP:

```
/dev/sct1/c6t110
```

On AIX:

```
/dev/ovpass0
```



That is all you need in the `mover.conf` file. In most cases, you can use the `bpmoverinfo` command to provide this entry.

You can use the following command to make sure the `sg` driver device path is correct. This command is also useful for creating the `mover.conf.policy_name` or `mover.conf.storage_unit_name` version of the `mover.conf` file (see [“Naming the Mover File”](#) on page 68).

```
/usr/opensv/volmgr/bin/sgscan
```

Here is some sample `sgscan` output showing third-party copy devices (see notes following):

```
/dev/sg/c0t16l10: Mover: "ADIC      Scalar SNC"/  
/dev/sg/c0t12l13: Mover: "ATTO      FB 4500H"  
/dev/sg/c0t15l10: Mover: "CNSi      FS2620"  
/dev/sg/c0t17l10: Mover: "PATHLIGHTSAN Gateway"  
/dev/sg/c0t11l13: Mover: "Crossrds4250 Router"  
/dev/sg/c0t21l10: Mover: "Crossrds8000 Router"  
/dev/sg/c0t23l12: Mover: "OVERLANDNEO VIA FCO"  
/dev/sg/c4t0l10: Changer: "SPECTRA 215"
```

Notes:

- ◆ “CNSi” indicates a Chaparral device.
- ◆ The number of entries returned for Crossroads depends on how many controller LUNS have been configured on that device. The `mover.conf` file must contain the `/dev` path for each controller LUN that is configured on the Crossroads.
- ◆ The Spectra Logic tape library does not have separate controller LUNs for the third-party functionality. For this reason, the `sgscan` output lists the library as a “Changer” rather than as a “Mover.”
- ◆ For an Hitachi, HP, or Sun disk array, you must check the HBA binding to obtain the SCSI target number for the array’s ECopy target port, and use that SCSI target number to identify the correct `/dev` path in the `sgscan` output.
- ◆ An alternative: the `mover.conf` file can consist of one line specifying the device by means of the `/dev/rmt/device_file_name` or `/dev/rdsk/device_file_name`, where `device_file_name` specifies the actual file name of the tape or disk. Note that the tape device must be the same as the device that NetBackup selects for the backup, and the disk must be one that is involved in the backup.
- ◆ Instead of the `/dev/rmt/device_file_name` or `/dev/rdsk/device_file_name` path, you can use the `TAPE` or `DISK` keyword. For more information, refer to [“Keywords in Mover File”](#) on page 64.

For sites that have multiple third-party copy devices

- ◆ In the `mover.conf` file, if you want to specify one of the third-party copy devices and prevent the others from being used, specify the device by means of its driver device path (such as `/dev/sg/c6t110` on Solaris, `/dev/sct1/c6t110` on HP, or `/dev/ovpass0` on AIX), or specify its `/dev/rmt/device_file_name` or `/dev/rdsk/device_file_name`.

See the list of example `/dev/sg` paths on the previous page.

- ◆ If you want to use any available tape or disk drive that is third-party-copy capable or connected to a third-party copy device, specify the `TAPE` or `DISK` keyword.
- ◆ If you want to limit the third-party copy device to that of a particular vendor or type, while including a variety of devices in the file, use the `END` keyword. First enter the device(s) you want to use, followed by `END`, then specify any other devices you might want to use at another time. For more information, refer to “[Keywords in Mover File](#)” on page 64.

SCSI Reserve/Release

For backups that use the third-party copy device method, SCSI reserve/release may be required to prevent unintended sharing of tape devices by multiple hosts on the SAN.

With SCSI reserve/release, either the media server or the third-party copy device acquires exclusive control over the tape drive, thus preventing other jobs from accessing the tape during the backup.

The `bptm` process logs all SCSI reserve/release commands. For background information on SCSI reserve/release, refer to the *NetBackup Media Manager System Administrator's Guide*.

To use SCSI reserve/release

Note If your `mover.conf` file contains only `/dev/rmt/device_path` entries or the `TAPE` keyword, SCSI reserve/release will be used for the backup. No further configuration is needed for SCSI reserve/release.

SCSI reserve/release is configured by means of the `mover.conf` file. The type of entry to make in `mover.conf` depends on the type of tape device and the network connection it is using, as follows:

- ◆ If the tape device is a Fibre Channel device (not connected behind a router or bridge) and does not have third-party copy functionality:

Specify the `passthru` path of the third-party copy device followed by the `i=reserve_value`. For example:



```
/dev/sg/c6t110 i=2000001086100d5e
```

where 2000001086100d5e is a 16-digit user-supplied value. See “[i=reserve_value](#)” on page 66 for more information on this value.

The third party copy device must be the tape device itself, or the tape drive must support a special kind of SCSI reserve/release called third-party reservation; otherwise, SCSI reserve/release will not be used. As of this writing, the only Advanced Client-supported device that supports third-party reservation is the ADIC/Pathlight Gateway. See “[i=reserve_value](#)” on page 66 for more information on the `reserve_value` required by the ADIC/Pathlight.

- ◆ If the tape device does not have third-party copy functionality and does not support the `i=reserve_value`:

Specify the passthru path of the third-party copy device followed by the `hr` keyword. For example:

```
/dev/sg/c6t110 hr
```

The `hr` keyword tells NetBackup to use SCSI reserve/release. If the `hr` keyword is omitted, SCSI reserve/release is not used.

- ◆ If the tape is behind the third-party copy device or has its own third-party copy functionality:

Specify the tape device path or the `TAPE` keyword. For example:

```
/dev/rmt/2cbn
```

or

```
TAPE
```

Keywords in Mover File

The following keywords can be included in the mover file:

DISK

For a third-party copy backup, attempt to use a disk involved with the current backup if that disk has third-party copy functionality or is behind (SCSI-connected to) a third-party copy device. This allows better concurrent backup processing, so that two or more backup jobs can execute simultaneously.

Note A valid SCSI passthru driver device path must be included in the `mover.conf` file, after the `DISK` keyword. This device path is used for tape verification and tape movement before a third-party copy capable disk is discovered. An example of a passthru driver device path is `/dev/sg/c6t110` on Solaris, `/dev/sctl/c6t110` on HP, or `/dev/ovpass0` on AIX.

TAPE

For a third-party copy backup, attempt to use the current tape device selected for the backup if that device has third-party copy functionality or is behind (SCSI-connected to) a third-party copy device. This has two advantages:

- ◆ There is no need to specify a device path or passthru driver device path. Instead of having to enter `/dev/rmt/` paths for a number of tape devices, you can use the `TAPE` keyword as shorthand for all of them.
- ◆ Allows better concurrent backup processing, so that two or more backup jobs can execute simultaneously.

GROUP *filename*

For a third-party copy backup, enable a third-party copy device to process two or more backup jobs simultaneously. This applies to devices that can handle multiple jobs simultaneously; not all third-party copy devices can do so. When enabled, simultaneous execution prevents multiple jobs waiting in a queue for the same device.

This keyword must be specified as `GROUP filename`, where *filename* is a file containing the device paths to be used for each simultaneous third-party copy backup. The file is assumed to be in the same directory as the `mover.conf` file, `/usr/opensv/volmgr/database`.

For example, the device paths for a third-party copy device that can run four jobs simultaneously might be as follows:

Solaris

```
/dev/sg/c0t010
/dev/sg/c0t011
/dev/sg/c0t012
/dev/sg/c0t013
```

HP

```
/dev/sctl/c6t110
/dev/sctl/c6t111
/dev/sctl/c6t112
/dev/sctl/c6t113
```



AIX

```
/dev/ovpass0  
/dev/ovpass1  
/dev/ovpass2  
/dev/ovpass3
```

END

Stop searching the `mover.conf` file for third-party copy devices for the current third-party copy backup.

If there are two or more third-party copy devices in the `mover.conf` file, NetBackup tries them sequentially, starting with the first one listed in the file, until one is found that can successfully move the data. `END` means do not look further in the current mover file and do not look in any other mover files, even if the last device tried was unsuccessful. Note that if no successful device is found before `END` is reached, the backup fails.

The `END` keyword limits the search for a third-party copy device in a `mover.conf` file that contains entries for more than one device. This can save you the trouble of deleting device entries and re-entering them later.

For example, if the `mover.conf` file contains the following:

```
/dev/sg/c6t4l0  
END  
/dev/sg/c6t4l2  
/dev/sg/c6t4l3
```

NetBackup will try to use device `/dev/sg/c6t4l0` and will not try the other devices.

The following optional keywords can be added to each entry in `mover.conf`

i=reserve_value

Use SCSI reserve/release for third-party reservation, if supported by the tape device or by the third-party copy device to which the tape device is connected. The `reserve_value` is a world-wide port name or fibre channel port identifier, as follows.

- ◆ For the ADIC/Pathlight Gateway, the `reserve_value` is the world-wide port name of the ADIC/Pathlight.
- ◆ For devices made by other vendors, the `reserve_value` may be the fibre channel port identifier (destination ID) of the third-party copy device, with two leading zeros. For example, if the fibre channel port identifier is 231DE4, the `reserve_value` is 00231DE4.

Please contact the vendor of the device for specifications.

hr

Hold the tape reservation (SCSI reserve/release) when a third-party copy device that is not a tape device is designated by means of a passthru device path (`/dev/sg/` on Solaris, `/dev/sct1/` on HP, `/dev/` on AIX). If you do not specify the `hr` keyword, the default is to drop or omit the reservation.

dr

Omit the use of SCSI reserve/release when a tape device is designated by the `TAPE` keyword or its tape device path (such as `/dev/rmt/2cbn`). If you do not specify the `dr` keyword, the default is to hold the reservation.

For example:

```
/dev/rmt/2cbn
/dev/rmt/3cbn
TAPE dr
```

In this example, if neither of the specified `/dev/rmt` devices can use SCSI reserve/release, NetBackup will try a tape device without the reserve.

to

If the third-party copy device needs additional time to respond to a backup request, you can increase the timeout value by specifying `to` followed by the limit in seconds. The default is 300 seconds (5 minutes). Additional time may be needed, for instance, if the third-party copy device is running in debug mode.

The following example resets the timeout for third-party copy device `/dev/rmt/2cbn` to 600 seconds:

```
/dev/rmt/2cbn to 600
```

In this example, NetBackup will allow the third-party copy device (accessible through `/dev/rmt/2cbn`) ten minutes to respond to a backup request. If the device does not respond within 10 minutes, NetBackup will try the next third-party copy device listed in the mover file. If no other devices are listed, the backup will fail.

A Note on Keywords for SCSI Reserve/Release

The same path (passthru or `/dev/rmt/device_name` path) can be specified several times with different keywords or no keywords. NetBackup tries each path in succession (whether or not they specify the same path), attempting to use SCSI reserve/release or not, as specified.

Example:

```
/dev/sg/c6t110 i=4873968475898744
```



```
/dev/sg/c6t110 hr  
/dev/sg/c6t110
```

In this example, NetBackup will try to use the third-party copy device specified by `/dev/sg/c6t110` and will attempt to use reserve/release by means of the `i=reserve_value`. If unsuccessful, NetBackup will try to use the same third-party copy device and reserve/release by means of the `hr` keyword (hold the reserve). If unsuccessful, NetBackup will use the third-party copy device without the reserve.

Naming the Mover File

In addition to the standard `mover.conf` file name, there are two other options for naming the mover file:

Per Policy

```
/usr/openv/volmgr/database/mover.conf.policy_name
```

where *policy_name* is the name of a NetBackup policy. All backups for this policy will use the third-party copy device specified in this `mover.conf.policy_name` file.

For a disk that has third-party copy device capability, use the `mover.conf.policy_name` to specify the disk as the third-party copy device for the policy that backs up that disk.

Per Storage Unit

```
/usr/openv/volmgr/database/mover.conf.storage_unit_name
```

where *storage_unit_name* is the name of a storage unit. This allows a third-party copy device to use a particular storage device by means of a storage unit name. Here is an example `mover.conf` file name of the *storage_unit_name* type:

```
mover.conf.nut-4mm-robot-tl4-0
```

where `nut-4mm-robot-tl4-0` was selected as the storage unit in the policy.

Note The *storage_unit_name* in this file name must exactly match the name of the storage unit as it appears in the “Policy storage unit” field of the Change Policy dialog.

Selection Priority for mover.conf files

NetBackup looks for an appropriate `mover.conf` file in the following order:

1. `mover.conf.policy_name`
2. `mover.conf.storage_unit_name`
3. `mover.conf`

Create the 3pc.conf File

The `/usr/opensv/volmgr/database/3pc.conf` file contains a list of all disk and tape devices on the SAN that NetBackup Advanced Client can use. NetBackup automatically creates this file at the start of the backup. In certain circumstances, however, you must create this file manually.

Note You must create a `3pc.conf` file if you are using the Third-Party Copy Device backup method AND some of your devices do not support identification descriptors (E4 target). Otherwise, you can skip to “[Create the mover.conf File](#)” on page 72.

1. Create a `3pc.conf` file as follows when no backups are in progress.

On the media server, run the `bptpcinfo` command:

```
/usr/opensv/netbackup/bin/bptpcinfo -a
```

- ◆ If a `3pc.conf` file already exists in `/usr/opensv/volmgr/database`, you are asked if you want to overwrite it. You can use the `-o output_file_name` option to send the output to a temporary file. However, the correct output must be included in the `/usr/opensv/volmgr/database/3pc.conf` file before backups can succeed.
- ◆ If a storage device is currently involved in a backup, the `bptpcinfo` command cannot gather information on that device and skips to the next device. If the `3pc.conf` file contains no entry for a storage device on your network, use the verbose mode (`-v`) of the `bptpcinfo` command to determine if the device was busy (see the *NetBackup Commands for UNIX* guide for more information on `bptpcinfo`).

2. If the media server does not have access to all disks (due to zoning or LUN-masking issues), run the following command on the media server:

```
/usr/opensv/netbackup/bin/bptpcinfo -x client_name
```

where *client_name* is the name of a NetBackup client on the fibre channel network where the third-party copy device is located. The `3pc.conf` file will be updated with information about the disks on this network, allowing the media server to “see” those disks. This information may have to be edited by adding the world-wide name (wwn=) of each device, as explained in the next two steps.



Note that the entries added by the `-x` option do not include `p=devpath`. Instead, they have `c=client` and `P=clientpath`. In the following example, lines 21 and 22 were added by the `-x` option:

```
16 p=/dev/rds/c4t1d3s2 s=FUJITSU:MAN3367MSUN36G:01X38525 l=3 a=100000E00221C153:3
17 p=/dev/rds/c4t1d4s2 s=FUJITSU:MAN3367MSUN36G:01X37951 l=4 a=100000E00221C153:4
18 p=/dev/rds/c4t1d5s2 s=FUJITSU:MAN3367MSUN36G:01X39217 l=5 a=100000E00221C153:5
19 p=/dev/rds/c4t3d1s2 s=HITACHI:OPEN-3-SUN:20461000300 l=1 i=10350060E800000000000004FED00000003 a=50060E80034FED00:1
20 p=/dev/rds/c4t3d2s2 s=HITACHI:OPEN-3-SUN:20461000400 l=2 i=10350060E800000000000004FED00000004 a=50060E80034FED00:2
21 s=SEAGATE:ST39204LCSUN9.0G:3BV0B09V0000104413ZF l=4 c=magner P=/dev/rds/c4t28d4s2
22 s=FUJITSU:MAG3182LSUN18G:01533134 l=8 c=magner P=/dev/rds/c4t28d8s2
```

3. If you have VERITAS CommandCentral Storage or SANPoint Control, you can use it to add world-wide name and lun information to the `3pc.conf` file, by entering the following command on the media server:

```
/usr/opensv/netbackup/bin/admincmd/bpSALinfo -S SPC_server
```

where `-S SPC_server` identifies the host where CommandCentral Storage (or SANPoint Control) is running. `bpSALinfo` adds world-wide name and lun values to device entries in the `3pc.conf` file. For additional command options, refer to the man page for `bpSALinfo`, or to the *NetBackup Commands for UNIX* guide.

Note If using CommandCentral Storage, the SAL remote component must be installed on all NetBackup clients in order for the `bpSALinfo` command to gather all required information.

If `bpSALinfo` successfully updates the `3pc.conf` file, no further editing of the `3pc.conf` file is required. You can skip to “[Create the mover.conf File](#)” on page 72.

4. If you do not have CommandCentral or SANPoint Control or it does not support your environment, edit the `3pc.conf` file as follows:

For each storage device listed in the `3pc.conf` file, you may need to provide world-wide port names, depending on what NetBackup was able to discover about the device and what the third-party copy device supports.

These are the editing tasks:

- ◆ In the `3pc.conf` file, if each device that will be backed up with Advanced Client has an identification descriptor (`i=`), and if the third-party copy device supports the use of identification descriptors, the `3pc.conf` file is complete. No editing is needed; skip the rest of this section and continue with “[Create the mover.conf File](#)” on page 72.

- ◆ If the `3pc.conf` file does not have an identification descriptor for each device (or the descriptor is not supported by the third-party copy device), enter the world-wide port name (w=) for each device. (Obtain the world-wide port name from your “[Device Checklist](#)” on page 52.)



Create the mover.conf File

This section describes how to create a list of the available third-party copy devices in a `mover.conf` file.

Note This is required for the Third-Party Copy Device method only.

1. On the NetBackup media server, enter the following command:

```
/usr/opensv/netbackup/bin/admincmd/bpmoverinfo
```

This creates the following file:

```
/usr/opensv/volmgr/database/mover.conf
```

The `bpmoverinfo` command discovers any third-party copy devices available on the SAN and lists them in the `mover.conf` file. Any tape drives with third-party copy capability are listed first.

For a description of the `bpmoverinfo` command, refer to the *NetBackup Commands for UNIX* guide.

Note For `bpmoverinfo` to correctly list third-party copy devices in the `mover.conf` file, the third-party copy devices must already have passthru paths defined. For an example, see “[Solaris only: Example for sg.links, sg.conf, and st.conf files](#)” on page 51.

2. If you need to control the circumstances under which a third-party copy device is used, create a separate `mover.conf` file for a policy or storage unit:

```
/usr/opensv/volmgr/database/mover.conf.policy_name
```

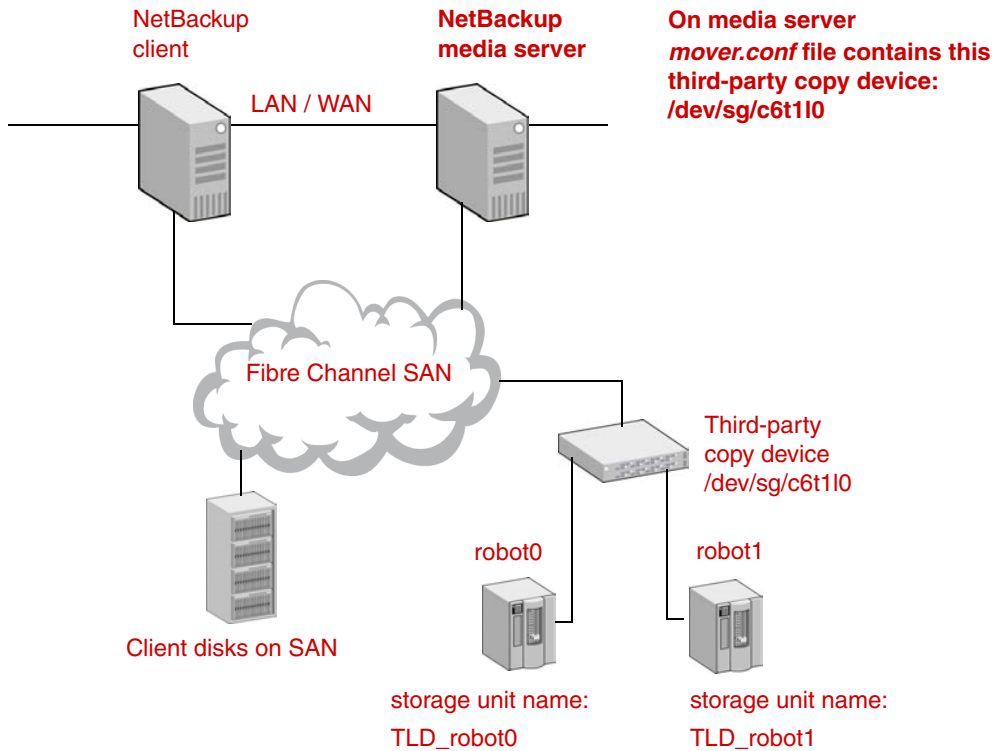
or

```
/usr/opensv/volmgr/database/mover.conf.storage_unit_name
```

For information on these naming formats and possible mover file entries, refer to “[mover.conf Description](#)” on page 59 and “[Naming the Mover File](#)” on page 68.

Following are example storage environments and `mover.conf` files.

Example mover.conf file for a site with one third-party copy device




In the above example, backups will use third-party copy device `/dev/sg/c6t110` specified in the `mover.conf` file. The backup uses the storage unit (TLD_robot0 or TLD_robot1) specified for the policy on the Change Policy dialog.

See the next figure for an example configuration involving a disk array with third-party copy device capability.



Example `mover.conf.policy_name` file for site with third-party copy capability in disk array



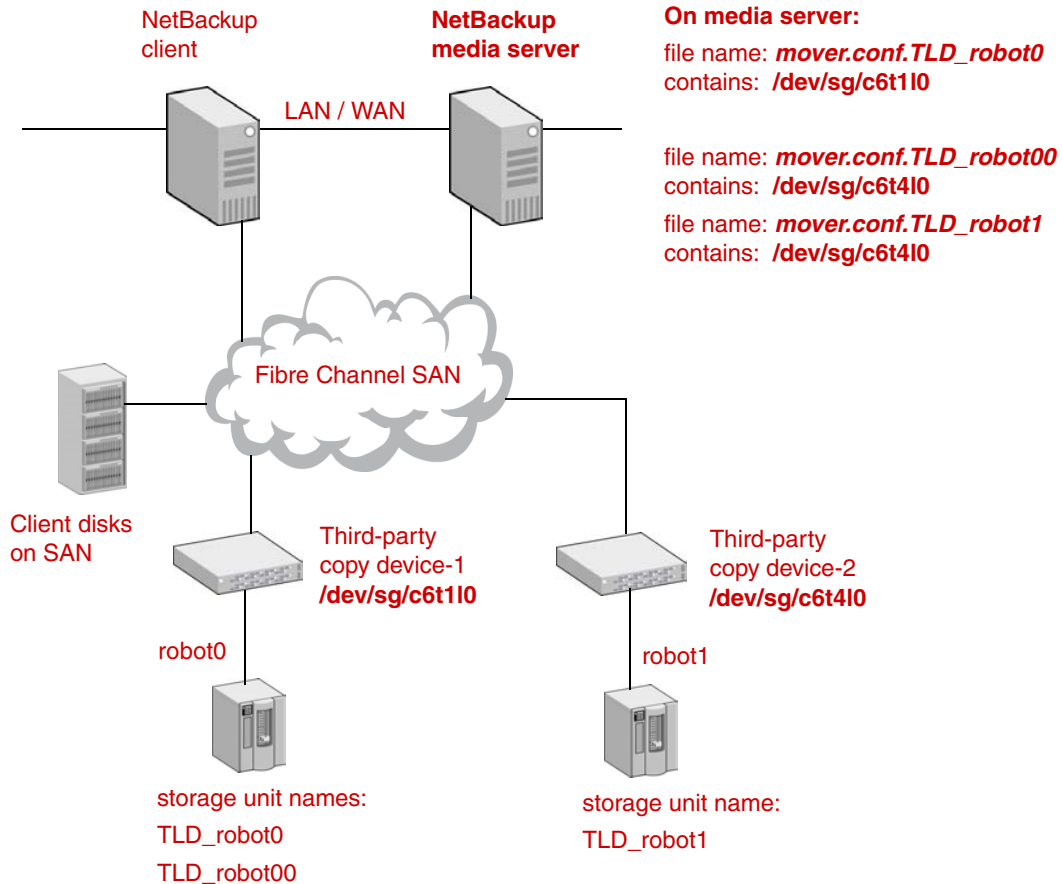
In this example, policy `array_1` is configured to back up the client data contained on the disk array. The backup uses storage unit `TLD_robot0` to store the data.

All backups configured in this policy will use the disk array as the third-party copy device. The `mover.conf.array_1` file specifies that array.

Note The client data must reside in the array that is used as the third-party copy device.

See the next figure for an example configuration with two third-party copy devices, where both devices can use the same robot.

Example mover.conf.storage_unit_name files for two third-party copy devices



The above example shows two robots (robot0 and robot1). Robot0 has been assigned two storage unit names, TLD_robot0 and TLD_robot00. Robot1 has been assigned one storage unit name, TLD_robot1.

The above example also shows two third-party copy devices, device-1 with a SCSI passthru device path of /dev/sg/c6t110, and device-2 with a SCSI passthru device path of /dev/sg/c6t410.

- ◆ To allow third-party copy device-1 to use robot0, create a file named mover.conf.TLD_robot0. In the file, include the device path of device-1 (/dev/sg/c6t110).



- ◆ To allow third-party copy device-2 to use the same robot (robot0), create a file named `mover.conf.TLD_robot00`. In the file, include the device path of device-2 (`/dev/sg/c6t4l0`). Notice that the file name must refer to a different storage unit, `TLD_robot00`, which is assigned to robot0.
- ◆ To allow third-party copy device-2 to use robot1, create a file named `mover.conf.TLD_robot1` that includes the device path of device-2 (`/dev/sg/c6t4l0`).

Note The *storage_unit_name* portion of the `mover.conf.storage_unit_name` file name must exactly match the actual name of the storage unit. See under [“Configuring an Advanced Client Policy”](#) on page 79 for an example Change Policy dialog showing a storage unit name in the **Policy storage unit** field.

Policy Configuration

4

This chapter is intended as detailed reference when using the Policies node of the NetBackup Administration Console to set up a policy for Advanced Client features.

The following topics are covered in this chapter:

- ◆ [Notes and Prerequisites](#)
- ◆ [Configuring an Advanced Client Policy](#)
- ◆ [Automatic Snapshot Selection](#)
- ◆ [Selecting the Snapshot Method](#)
- ◆ [Configuring Backup Scripts](#)
- ◆ [Configuring Alternate Client Backup](#)
- ◆ [Configuration Tips](#)



Notes and Prerequisites

Before proceeding with this chapter, note the following.

- ◆ NetBackup Enterprise and the Advanced Client add-on product must be installed on master server(s) and clients.
- ◆ For the NetBackup Media Server and Third-Party Copy Device offhost backup methods, a fibre channel network or multi-ported SCSI disk array must be configured.
- ◆ Storage devices must be configured (you can use the Device Configuration wizard).
- ◆ To use BLIB with Advanced Client (**Perform block level incremental backups** option on the policy attributes tab), you must install the NetBackup for Oracle or NetBackup for DB2 software. Oracle and DB2 are the only client data types supported by the block level incremental backup feature of Advanced Client. Refer to the appropriate NetBackup database agent guide for assistance in making block level incremental backups with Advanced Client.

Note If you choose the **Perform block level incremental backups** option on the policy attributes tab, the other features of Advanced Client are not available and are grayed out.



Configuring an Advanced Client Policy

As an alternative to this procedure, you can use the **Snapshot Policy Configuration wizard**.

1. Start the NetBackup Administration Console as follows:
On UNIX, enter: `/usr/openv/netbackup/bin/jnbSA &`
On Windows, click **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**.
2. Click on **Policies**. In the **All Policies** pane, double click on the name of the policy (or right-click to create a new one).

Change Policy – sample

Server: be

Attributes Schedules Clients Backup Selections

Policy type: Standard

Destination:

Policy storage unit: Any available

Policy volume pool: NetBackup

☐ Take checkpoints every: 0 minutes

☐ Limit jobs per policy: 0

Job priority: 0 (higher number is greater priority)

☒ **Active. Go into effect at:** 06/30/2005 18:15:25

☐ Follow NFS

☐ Cross mount points

☐ Compression

☐ Encryption

☐ Collect disaster recovery information for Intelligent Disaster

☐ Collect disaster recovery information for Bare Metal Restore

☐ Collect true image restore information

☐ with move detection (Required for synthetic backups and Bare Metal Restore)

☐ Allow multiple data streams

Keyword phrase (optional):

Advanced Client

☐ Perform block level incremental backups

☒ **Perform snapshot backups** [Advanced Snapshot Options...](#)

☐ Retain snapshots for Instant Recovery

☒ **Perform offhost backup**

☒ Use alternate client ☐ Use data mover

Buttons: Apply OK Close Help

Annotations:

- Select the policy type.
- Select appropriate storage unit or storage unit group.
- Requires NetBackup Oracle or DB2 database agent software. See the appropriate NetBackup database agent guide.
- Click **Perform snapshot backups**.
- If you checked **Perform Offhost backup**, specify method.



3. Select the policy type:

- ◆ If client data is in a database, select the database type (**DB2**, **Oracle**, **MS-Exchange-Server**, **MS-SQL-Server**).
- ◆ To use FlashBackup, select **FlashBackup** for UNIX clients or **FlashBackup-Windows** for Windows clients. See the “[FlashBackup Configuration](#)” chapter for further information.
- ◆ For all other cases, select **Standard** for UNIX clients and **MS-Windows-NT** for Windows clients.

4. Make sure **Perform snapshot backups** is selected.

If Bare Metal Restore is installed, you must deselect **Collect disaster recovery information for Bare Metal Restore** in order to select **Perform snapshot backups**.

5. Optional: To select the snapshot method manually, refer to “[Selecting the Snapshot Method](#)” on page 86. Skip this step if you want NetBackup to select the snapshot method for you. “[Automatic Snapshot Selection](#)” on page 84 describes how NetBackup chooses a snapshot method.

6. To create a backup that enables Instant Recovery, select the **Retain snapshots for instant recovery** attribute.

This attribute is required for the following types of restore: block-level restore, file promotion, and rollback. These are described under “[Instant Recovery Restore](#)” on page 196. For help in creating a policy for instant recovery backups, refer to “[Instant Recovery Configuration](#)” on page 123.

7. To reduce the processing load on the client, select **Perform offhost backup**.

This option may require additional configuration.

- a. For a backup performed by an alternate client, select **Use alternate client** and enter the name of the alternate client. Refer to “[Configuring Alternate Client Backup](#)” on page 93 for more information.
- b. For a backup performed by a data mover (not by a client), select **Use data mover** and select the method:

NetBackup Media Server

Backup processing will be handled by a Solaris, HP, AIX, or Linux NetBackup media server (for Solaris, HP, AIX, and Linux clients only).

Third-Party Copy Device

Backup processing will be handled by a third-party copy device (for Solaris, HP, AIX, and Linux clients only).



Network Attached Storage

Backup processing will be handled by an NDMP host (NAS filer), with the NAS_Snapshot method. NetBackup for NDMP software is required. For help configuring a policy for Network Attached Storage and NAS_Snapshot, refer to the [“NAS Snapshot Configuration”](#) chapter.

- c. Specify a policy storage unit or group of storage units in the **Policy storage unit** pull-down menu.

Note **Any_available** is not supported for the following data mover types: NetBackup Media Server and Third-Party Copy Device. Disk storage units are not supported for Third-Party Copy Device.

Instead of using a particular storage unit, you can create a storage unit group that designates devices configured on the SAN. Storage unit groups are described in the *NetBackup Media Manager System Administrator's Guide*.

- d. Also note the following:
 - ◆ For the Third-Party Copy Device option, you must create a mover.conf file. See [“Create the mover.conf File”](#) on page 72. For background information on this file, see [“mover.conf Description”](#) on page 59.
 - ◆ If you do not have VERITAS CommandCentral Storage (or SANPoint Control), and one or more devices needed for the backup do not support identification descriptors (SCSI E4 target descriptors), you may need to edit the 3pc.conf file as explained in [“Create the 3pc.conf File”](#) on page 69.

8. To save these settings, click **Apply**.
9. Use the **Schedules** tab to define a schedule, and the **Clients** tab to specify the clients.

Regarding clients: only one snapshot method can be configured per policy. If you want to select one snapshot method for clients a, b, and c, and a different method for clients d, e, and f, create a separate policy for each group of clients and select one method for each policy. (You may be able to avoid this restriction using the auto method; see [“The auto method has these advantages:”](#) on page 84.)

10. Use the **Backup Selections** tab to specify the files to be backed up. Please note:
 - ◆ Advanced Client policies do not support the ALL_LOCAL_DRIVES entry in the Backup Selections list.



- ◆ For snapshot backups, the maximum pathname length is approximately 1000 characters (as opposed to 1023 characters for backups that do not use a snapshot method). Refer to “[Maximum Pathname Length](#)” on page 96 for more information on this restriction. The *NetBackup System Administrator’s Guide* describes other file-path rules.
- ◆ Wildcards are permitted, provided that the wildcard does not correspond to a mount point or is not followed by a mount point in the path.

For example, in the path `/a/b`, if `/a` is a mounted file system or volume, and `/a/b` designates a subdirectory in that file system, then an entry of `/a/b/*.pdf` will cause NetBackup to make a snapshot of the `/a` file system and to back up all pdf files in the `/a/b` directory. But, with an entry of `/*` or `*/b`, the backup may fail or have unpredictable results, because the wildcard corresponds to the mount point `/a`. Do not use a wildcard to represent all or part of a mount point.

In a second example, `/a` is a mounted file system which contains another mounted file system at `/a/b/c` (where `c` designates a second mount point). A Backup Selections entry of `/a/*/c` may fail or have unpredictable results, because the wildcard is followed by a mount point in the path.

See also “[Snapshot Tips](#)” on page 96 for a note about the **Cross mount points** policy attribute.

- ◆ For a raw partition backup of a UNIX client, specify the `/rdsk` path, not the `/dsk` path. You can specify the disk partition (except on AIX or Linux) or a VxVM volume.

Examples:

On Solaris:

```
/dev/rdsk/c0t0d0s1  
/dev/vx/rdsk/volgrp1/vol1
```

On HP:

```
/dev/rdsk/c1t0d0  
/dev/vx/rdsk/volgrp1/vol1
```

On AIX and Linux:

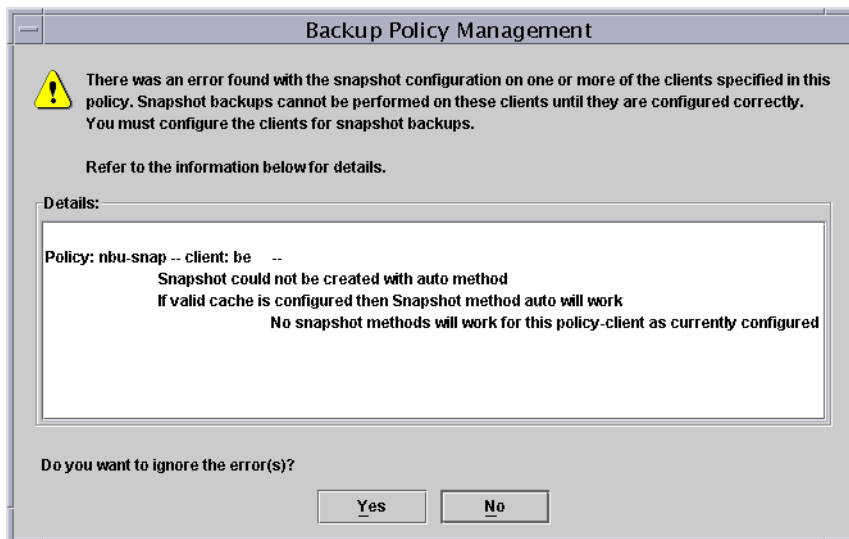
```
/dev/vx/rdsk/volgrp1/vol1
```

Note On AIX and Linux clients, backing up a native disk partition is not supported. A raw partition backup must specify a VxVM volume, such as `/dev/vx/rdsk/volgrp1/vol1`. Note that `/dev/vx/dsk/volgrp1/vol1` (without the “r” in `/rdsk`) will not work.

- ◆ If you use the Backup Policy Configuration wizard, see “[Backup Policy Configuration Wizard](#)” on page 97.

11. Click *Close* when done with Schedules and Clients tabs.

Advanced Client validation begins. An error message may be displayed, such as the following:



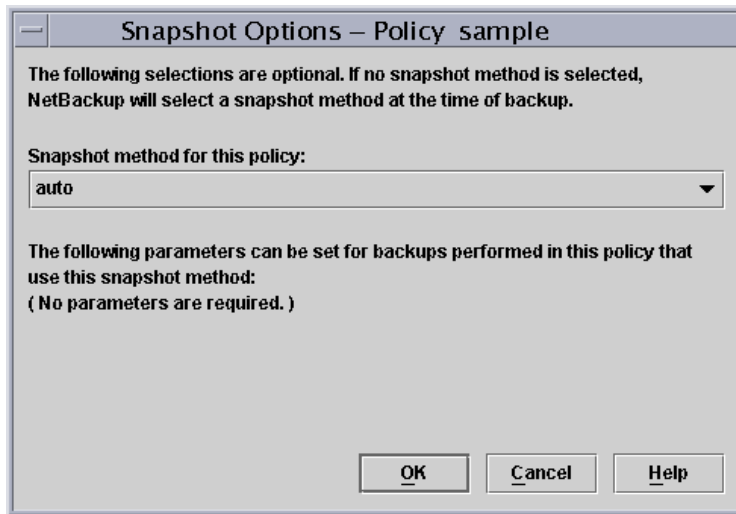
The Details pane explains the problem. You can click **No**, resolve the problem, and close the policy again, or click **Yes** to override the message.



Automatic Snapshot Selection

To have NetBackup select the snapshot method, click **Perform snapshot backups** on the policy Attributes tab.

- ◆ If this is a new policy, NetBackup will select a snapshot method when the backup starts (by default, the snapshot method is set to **auto**).
- ◆ If this is a copy of a policy that was configured for a particular snapshot method, click the **Advanced Snapshot Options** button and set the snapshot method to **auto**. NetBackup will select a snapshot method when the backup starts.



Use of the **auto** method does not guarantee that NetBackup can select a snapshot method for the backup. NetBackup looks for a suitable method based on a number of factors:

- ◆ The client platform and policy type.
- ◆ The presence of up-to-date software licenses, such as VERITAS VxFS and VxVM.
- ◆ How the client data is configured. For instance, whether a raw partition has been specified for a copy-on-write cache (refer to [“How to Enter the Cache”](#) on page 145), or whether the client’s data is contained in VxVM volumes configured with one or more snapshot mirrors.

NetBackup uses the first suitable method found.

The auto method has these advantages:

- ◆ It allows NetBackup to use a different snapshot method for each item in the Backup Selections list, or for each client in the policy, if need be. This gives NetBackup more flexibility in choosing a snapshot method and allows you to circumvent the one-snapshot-method-per-policy restriction.

- ◆ At the time of backup, NetBackup selects a snapshot method that will work based on how the client is configured at that moment. If, for instance, the VxFS or VxVM license on the client has expired, or if the client has been recently reconfigured, NetBackup takes this into account in selecting a snapshot method.



Selecting the Snapshot Method

This section explains how to manually configure a snapshot method from the Administration Console.

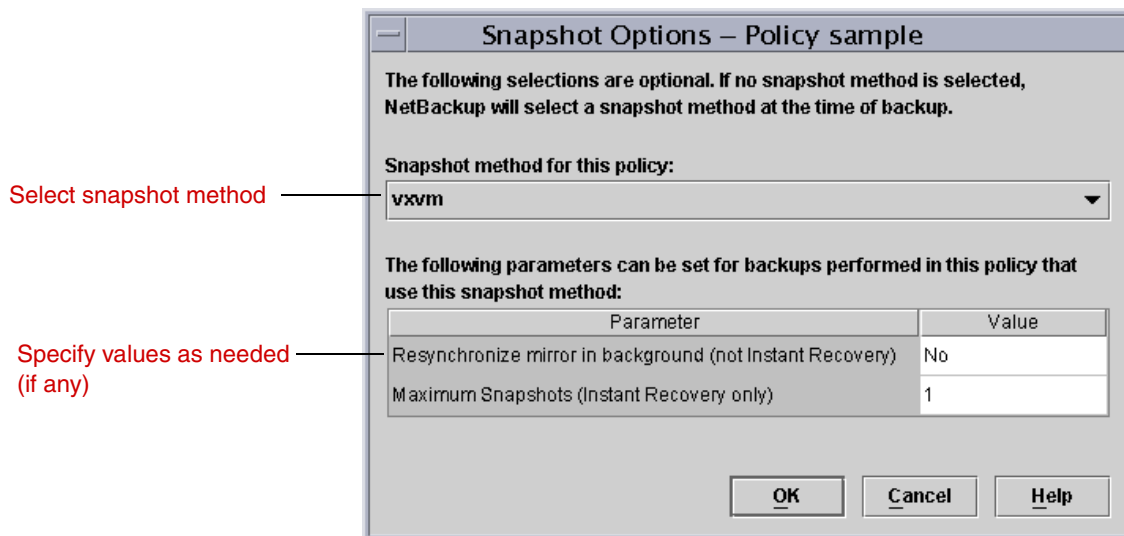
1. Start the NetBackup Administration Console:

On UNIX, enter: `/usr/openv/netbackup/bin/jnbSA &`

On Windows, click **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**.

2. Click on **Policies**. In the **All Policies** pane, double click on the name of the policy. The Change Policy dialog appears.
3. Make sure that **Perform snapshot backups** is selected (for more information on this, refer to the previous procedure).
4. Click **Advanced Snapshot Options**.

The Snapshot Options dialog appears. This is for selecting a snapshot method for the policy.



5. Select the snapshot method for the policy.

The available methods depend on how your clients are configured and which attributes you selected on the Attributes tab.

Note Only one snapshot method can be configured per policy. If you want to select one snapshot method for clients a, b, and c, and a different method for clients d, e, and f, create two policies for each group of clients and select one method for each policy.

Note The snapshot method you select must be compatible with all items in the Backup Selections list.

auto

NetBackup will select a snapshot method when the backup starts. See “[Automatic Snapshot Selection](#)” on page 84. If necessary, NetBackup will select a different method for each item in the Backup Selections list.

BusinessCopy

For mirror snapshots with Hewlett Packard disk arrays with BusinessCopy Services. For clients on Solaris or HP.

FlashSnap

For mirror snapshots on alternate clients, with the VxVM FlashSnap feature. This method is for clients on Solaris, HP, AIX, Linux, and Windows. UNIX clients must be at VxVM 3.2 or later; Linux and AIX clients at VxVM 4.0 or later; Windows clients must be at VxVM 3.1 or later, with all the latest VxVM service packs and updates.

FlashSnap is based on the VxVM disk group split and join technology.

NAS_Snapshot

For copy-on-write snapshots of client data residing on an NDMP host. Requires NetBackup for NDMP software. For help configuring a policy for NAS_Snapshot, refer to the “[NAS Snapshot Configuration](#)” chapter.

nbu_snap

For copy-on-write snapshots of UFS or VERITAS VxFS file systems. For Solaris clients only. nbu_snap is not supported in clustered file systems.

nbu_snap requires a designated cache; see “[How to Enter the Cache](#)” on page 145.

ShadowImage

For mirror snapshots with Hitachi Data Systems disk arrays with ShadowImage (HOMRCF). For clients on Solaris, HP, Linux, and AIX.

TimeFinder

For mirror snapshots with EMC disk arrays (with TimeFinder SYMCLI). For clients on Solaris, HP, Linux, and AIX.



VSP

VERITAS Volume Snapshot Provider, for snapshots of open and active files. For Windows clients only.

Note You can use VSP without Advanced Client, as explained in the *NetBackup System Administrator's Guide for Windows, Volume I*. In some cases, however, such as when the **Busy File Timeout** has expired, no snapshot is created and the backup job may continue without backing up the busy file. If you use VSP with Advanced Client, the backup will either successfully create a snapshot of all files, or the backup job will fail.

VSS

For snapshots using the Volume Shadow Copy Service of Windows 2003. This method is for clients on Windows 2003 only. Instant Recovery is not supported. The note under VSP also applies to VSS.

Note VSS is a general interface to Windows Shadow Copy Services. The actual snapshot method used for the backup depends on the snapshot providers available on your system. When VSS is configured in the NetBackup policy, Windows Shadow Copy Services selects the actual snapshot mechanism.

For configuration assistance, please refer to your Microsoft documentation.

VSS_Transportable

For snapshots using the Volume Shadow Copy Service of Windows 2003. This method is for alternate client backup of Windows 2003 clients, where the client data is stored on a disk array or in a VERITAS Storage Foundation for Windows 4.1 or later volume. It supports file system backup of a disk partition (such as E:\) and backup of Exchange databases. The disk partition must be on a disk array that has been configured with disk pairing and masking as explained under “[VSS_Transportable](#)” on page 159.

For a list of supported disk arrays, refer to “[Advanced Client Assistance](#)” on page xvi.

VVR

For alternate client backups of a replicated VxVM volume. For clients on Solaris, HP, Linux, AIX.

Requires VxVM 3.2 or later with the VERITAS Volume Replicator license. Linux and AIX clients require VxVM 4.0 or later.

VxFS_Checkpoint

For copy-on-write snapshots of clients on Solaris, HP, AIX, or Linux. This method is not supported by the FlashBackup policy type.

Requires the Storage Checkpoint feature of VxFS 3.4 or later. HP requires VxFS 3.5 or later. Linux and AIX clients require VxFS 4.0 or later.

Note Note that VxFS_Checkpoint requires the NetBackup Advanced Client license and the VERITAS File System license with the Storage Checkpoints feature. Without both licenses, the copy-on-write snapshot (Storage Checkpoint) cannot be opened and the backup fails.

VxFS_Snapshot

For snapshots of Solaris or HP clients on the local host (not offhost), for FlashBackup policies only. This method requires VxFS 3.4 (Solaris) or VxFS 3.3.2 (HP) or later. This method also requires a designated cache; see [“VxFS_Snapshot”](#) on page 149 for details. Note that all files in the Backup Selections list must reside in the same file system.

vxvm

For any of the following types of snapshots with data configured over Volume Manager volumes, for clients on Solaris, HP, AIX, Linux, or Windows. (Linux and AIX clients require VxVM 4.0 or later.)

- ◆ For “third-mirror” snapshots (VxVM 3.1 or later).
- ◆ For full-sized instant snapshots (VxVM 4.0).
- ◆ For space-optimized instant snapshots (VxVM 4.0).

Note For further notes relating to these snapshots, refer to the [“Instant Recovery Configuration”](#) chapter and the [“Snapshot Configuration Notes”](#) chapter.

6. Specify required parameters (if any).

Cache device path:

Specify a raw partition for the cache (as either a logical volume or a physical disk) by entering the cache partition’s full path name in the **Value** field. For example:

Solaris raw partition:

```
/dev/rdsk/c2t0d3s3
```

VxVM volume:

```
/dev/vx/rdsk/diskgroup_1/volume_3
```

HP LVM volume:

```
/dev/volume_group_1/volume_3
```



This setting overrides the cache specified on **Host Properties > Clients > Client Properties dialog > UNIX Client > Client Settings** (see “[How to Enter the Cache](#)” on page 145).

Do not specify wildcards (such as `/dev/rdisk/c2*`). See “[Cache device](#)” on page 142 for a complete list of requirements.

Caution The cache partition’s contents will be overwritten by the `nbu_snap` or `VxFS_Snapshot` process.

Keep snapshot after backup:

This option specifies whether or not the snapshot image is retained on the mirror disk after the backup completes (default is **No**). Retaining the image (**Yes**) enables a quick restore from the mirror in case data is deleted from the client’s primary disk. The image is retained on the mirror until the next time a backup is run using this policy. **No** indicates that the image will be deleted from the mirror as soon as the backup completes.

Note If the client is rebooted, snapshots that have been kept must be remounted before they can be used for a restore. You can use the `bpfis` command to discover the images (refer to the `bpfis` man page or the *NetBackup Commands* manual). This does not apply to snapshots for Instant Recovery: NetBackup automatically remounts them as needed.

Note If the snapshot is made on an EMC, Hitachi, or HP disk array, and you want to use hardware-level restore, read the **Caution** under “[Hardware-Level Disk Restore](#)” on page 186.

Maximum Snapshots (Instant Recovery only):

This is one of two options for determining when Instant Recovery snapshots are deleted. (The other option is backup retention level, which is set under the **Host Properties > Master Server Properties** dialog.) Refer to “[Configuring Snapshot Deletion](#)” on page 128 for an explanation of these two options.

Caution If you specify a number for maximum snapshots that is smaller than the existing number of snapshots, NetBackup deletes the older snapshots until the number of snapshots equals that specified for **Maximum Snapshots**.

Synchronize mirror before backup:

At start of the backup, this option determines whether or not the primary and mirror devices are automatically synchronized (if they were not already synchronized) before the backup begins. Default is **No**. Specify **Yes** to have unsynchronized devices

synchronized before the backup begins. **No** means that unsynchronized devices will not be synchronized before the backup starts. In this case (**No**), if the primary and mirror devices are not synchronized, the backup will fail.

Resynchronize mirror in background:

At end of the backup, this option determines whether or not the backup job will complete before the resynchronize operation has finished (default is **No**). **Yes** means that a resynchronize request will be issued, but the backup will not wait for the resync operation to complete. **No** means that the backup job cannot reach completion until the resynchronize operation has finished.

Choosing **Yes** may allow more efficient use of backup resources. For example, if two backup jobs need the same tape drive, the second job can start even though the resynchronize in background operation issued by the first job has not completed.

IBC send timeout (seconds):

For the VVR snapshot method, this value determines how long NetBackup waits for the next end-of-transaction marker to be automatically removed from the replication data stream. If the marker is not removed, replication cannot continue. If this timeout is exceeded before the current marker is removed, the backup will fail.

IBC receive timeout (seconds):

For the VVR snapshot method, this value determines how long NetBackup waits for the next end-of-transaction marker to be received in the replication data stream. For instance, a process may fail to send the marker, or the network connection may be down or overloaded. If this timeout is exceeded before the next marker is received, the backup will fail.

7. When finished, click **OK**.



Configuring Backup Scripts

For backups using a snapshot method, you can run scripts before and after the snapshot by adding directives to the Backup Selections list, as follows.

1. Start the NetBackup Administration Console by entering the following:

On UNIX: `/usr/opensv/netbackup/bin/jnbSA &`

On Windows: **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**

2. Click on **Policies**. In the **All Policies** pane, double click on the name of the policy. The Change Policy dialog appears.
3. Click the **Backup Selections** tab.
4. Add the following directive to the start of the Backup Selections list:

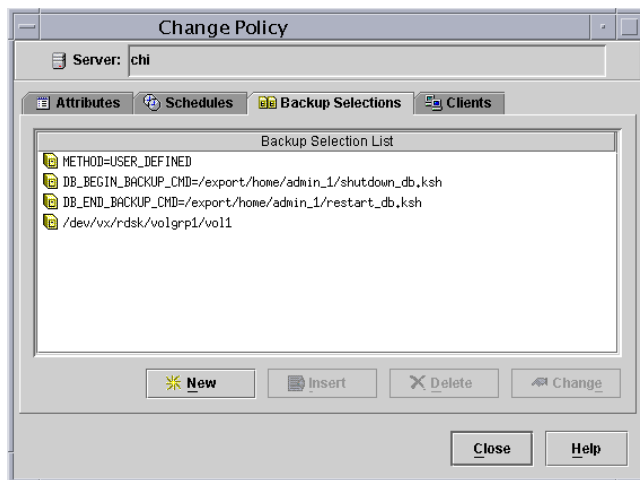
```
METHOD=USER_DEFINED
```

5. Add one or both of the following directive(s), as required.

```
DB_BEGIN_BACKUP_CMD=your_begin_script_path
DB_END_BACKUP_CMD=your_end_script_path
```

The arguments (arg) are optional.

For example:



In the above example, the script `shutdown_db.ksh` is run before the backup, and `restart_db.ksh` is run after the snapshot is created.

Configuring Alternate Client Backup

NetBackup Advanced Client supports three methods of doing offhost backup: alternate client backup, NetBackup Media Server, and Third-Party Copy Device. This section discusses alternate client backup.

Note that alternate client backup does not require a SAN, a third-party copy device, or Fibre Channel equipment. There is no need to create a `3pc.conf` or `mover.conf` file, or to do any of the procedures described in the [“SAN Configuration for Advanced Client”](#) chapter.

Alternate client backup is offhost in that all backup processing is off-loaded to another client. Off-loading the work to another client saves computing resources on the original client. The backup I/O processing is handled by the alternate client, and the backup has little impact on the original client. For an introduction to the basic configurations supported for alternate client backup (split mirror and data replication), refer to [“Alternate Client Backup”](#) on page 10.

Basic Requirements

Before configuring a policy for alternate client backup, make sure the following have been done:

- ◆ The client data must be available as either split mirrors or replication. The alternate client must have access to the mirror disk(s) or replication disk(s). See [“Data Sharing Between Clients”](#) on page 11 for more details on configuration.
- ◆ For the FlashSnap and VVR snapshot methods, VxVM 3.2 or later for UNIX, VxVM 4.0 or later for Linux and AIX, or VxVM 3.1 or later for Windows must be installed and volumes configured over the primary host’s disks. The VxVM FlashSnap or VVR license must also be installed.
- ◆ The user and group identification numbers (UIDs and GIDs) associated with the files to be backed up must be available to both hosts (the primary client and the alternate client).
- ◆ The primary and alternate clients must be running the same operating system, volume manager, and file system. For each of these I/O system components, the alternate client must be at the same level as the primary client, or higher level. Following are the supported configurations:

If primary client is:	Alternate client must be:
Windows	Windows, at same level as primary client or higher
Solaris	Solaris, at same level as primary client or higher



If primary client is:	Alternate client must be:
HP	HP, at same level as primary client or higher
AIX	AIX, at same level as primary client or higher
Red Hat	Red Hat, at same level as primary client or higher
SUSE	SUSE, at same level as primary client or higher
VxFS 3.4 or later (VxFS 3.3 for HP)	VxFS, at same level as primary client or higher
VxVM 3.2 or later (UNIX) VxVM 3.1 or later ¹ (Windows)	VxVM, at same level as primary client or higher. Note: for VVR method, the alternate client must be at exactly the same level as primary client.

1. For VxVM on Windows, use VxVM 3.1 or later with all the latest VxVM service packs and updates. See [“Available Backup Methods and Snapshot Methods”](#) for further requirements.

Note Alternate client backup on Windows does not support incremental backups based on archive bit. Instead, use incremental backups based on timestamp. See [“Incremental Backup of Mirror-Based Snapshots”](#) on page 98 for more information.

Available Backup Methods and Snapshot Methods

When setting up a policy for alternate client backup, you can make the following selections for policy type, offhost backup method, and snapshot method, depending on your hardware configuration and NetBackup add-on product licenses:

- ◆ **Policy type:** Choose Standard, FlashBackup, FlashBackup-Windows, MS-Windows-NT, MS-Exchange-Server, MS-SQL-Server, DB2, or Oracle.
- ◆ Click **Perform snapshot backups**.
- ◆ Click **Perform offhost backup**.
- ◆ Click **Use alternate client**, and select the alternate client from the drop-down list, or type it in.
- ◆ Click **Advanced Snapshot Options** if you want to choose the snapshot method.

Snapshot method: You can select the **auto** method, or the following:

- ◆ FlashSnap, for a disk group split configuration, with VxVM 3.2 or later with the FlashSnap feature.

- ◆ VVR, for a UNIX replication host; requires VxVM 3.2 or later with VVR feature.
- ◆ VSS_Transportable, for snapshots using the Volume Shadow Copy Service of Windows 2003. This method is for Windows 2003 clients, where the client data is stored on a disk array such as EMC or Hitachi, or in a VERITAS Storage Foundation for Windows 4.1 or later volume. Supports Exchange.
- ◆ TimeFinder, ShadowImage, BusinessCopy (the array-related methods, UNIX only).

Example configurations

1. Client data is on an EMC disk array in split-mirror mode

To run the backup on an alternate client, choose **Standard** as the policy type, select **Perform snapshot backups**, **Perform offhost backup**, **Use alternate client**, and select the alternate client. On the Snapshot Options display, specify the **TimeFinder** snapshot method.

If the data is in an Oracle database, select **Oracle** as the policy type.

2. Client data is replicated on a remote host

To run the backup on the replication host (alternate client), choose **Standard** as the policy type, select **Perform snapshot backups**, **Perform offhost backup**, **Use alternate client**, and select the alternate client (the replication host). On the Snapshot Options display, specify the **VVR** snapshot method.

3. Client data is on a JBOD array in VxVM volumes with snapshot mirrors configured

To run the backup on the alternate client, choose **Standard** (for UNIX client) or **MS-Windows-NT** (Windows client) as the policy type, select **Perform snapshot backups**, **Perform offhost backup**, **Use alternate client**, and select the alternate client. On the Snapshot Options display, specify the **FlashSnap** method.

If the client data consists of many files, or if you need the ability to do individual file restore from raw partition backups, select **FlashBackup** or **FlashBackup-Windows** as the policy type.

Note Other combinations of policy type and snapshot method are possible, depending on many factors, such as your hardware configuration, file system and volume manager configuration, and installed NetBackup product licenses.



Before running the alternate client backup

Your volume configuration must be prepared and tested for the snapshot method you will use (see the “[Snapshot Configuration Notes](#)” chapter).

Configuration Tips

The following items relate to policy creation.

Maximum Pathname Length

For snapshot backups, the maximum file list pathname length is approximately 1000 characters (as opposed to 1023 characters for backups that do not use a snapshot method). This is because the snapshot is created on a new mount point which is derived from the original Backup Selections list pathname. If this new mount point plus the original file path exceeds the system-defined maximum path name length (1023 characters on many systems), the backup fails with a status code 1, “the requested operation was partially successful.”

Refer to the *NetBackup System Administrator’s Guide* for other NetBackup file-path rules.

Snapshot Tips

- ◆ In the Backup Selections list, be sure to specify absolute path names. Refer to the *NetBackup System Administrator’s Guide* for help specifying files in the Backup Selections list.
- ◆ If an entry in the Backup Selections list is a symbolic (soft) link to another file, Advanced Client backs up the *link*, not the file to which the link points. This is standard NetBackup behavior. To back up the actual data, include the file path to the actual data.
- ◆ On the other hand, a raw partition can be specified in its usual symbolic-link form (such as `/dev/rdsk/c0t1d0s1`): do not specify the actual device name that `/dev/rdsk/c0t1d0s1` is pointing to. For raw partitions, Advanced Client automatically resolves the symbolic link to the actual device.
- ◆ The **Cross mount points** policy attribute is not available for policies that are configured for snapshots. This means that NetBackup will not cross file system boundaries during a backup of a snapshot. A backup of a high-level file system, such as `/` (root), will not back up files residing in lower-level file systems unless those file systems are also specified as separate entries in the Backup Selections list. To back up `/usr` and `/var`, for instance, both `/usr` and `/var` must be included as separate entries in the Backup Selections list.



For more information on **Cross mount points**, refer to the *NetBackup System Administrator's Guide*.

- ◆ On Windows, the \ must be entered in the Backup Selections list after the drive letter (for example, D:\). For the correct format when using a FlashBackup-Windows policy, see [step 10](#) on page 103.

Backup Policy Configuration Wizard

If you used the Backup Policy Configuration wizard to configure a policy for NetBackup, please note the following:

- ◆ If you chose **Back up all local drives**, the wizard placed an entry called ALL_LOCAL_DRIVES in the policy Backup Selections list. Advanced Client does not support the ALL_LOCAL_DRIVES option. Be sure to remove the ALL_LOCAL_DRIVES entry from the Backup Selections list.
- ◆ The wizard automatically sets your storage device to **Any_available**. This is correct for local backups. However, to use the Third-Party Copy Device or NetBackup Media Server method, you must select a particular storage unit—do not select **Any_available**.

Note To create an Advanced Client policy by means of a wizard, use the Snapshot Policy Configuration Wizard, which is specifically designed for Advanced Client features.

Multiple Data Streams

Note Multiplexing is not supported for the Third-Party Copy Device method. (Keep in mind that multiplexing is not the same as multiple data streams.)

For backups, make sure the following are set to allow the number of active streams to be equal to or greater than the number of streams in the Backup Selections list:

- ◆ Policy attribute: **Limit jobs per policy**
- ◆ Schedule setting: **Media multiplexing**
- ◆ Storage unit setting: **Maximum multiplexing per drive**
- ◆ System configuration setting: **Maximum jobs per client**



Incremental Backup of Mirror-Based Snapshots

For incremental backup of a mirror type snapshot, either local or remote (offhost) backup, note the following issues.

Archive Bit Incrementals (Windows only)

NetBackup provides two incremental backup options for Windows clients:

- ◆ Based on timestamps
- ◆ Based on archive bit (for Windows clients only)

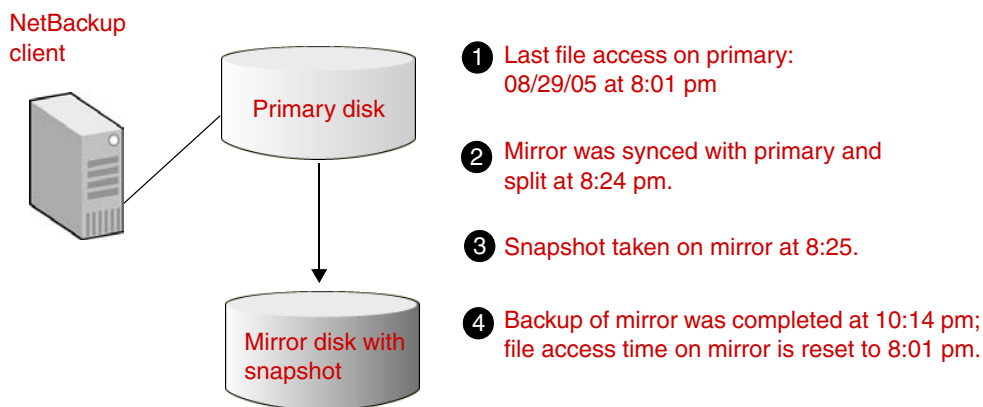
These options are available in the NetBackup Administration Console under **Host Properties** for the client, under **Windows Client > Client Settings**, and are described in the *NetBackup System Administrator's Guide, vol I*.

If the snapshot is created on a mirror, such as a Volume Manager volume using the vxvm method, incremental backups based on archive bit cannot be used. The archive bit incrementals are incompatible with snapshots on mirrors. For incrementals, select Based on timestamps.

Access Time Not Updated On Mirror After Backup

After NetBackup backs up a mirror snapshot, the access time of the files in the mirror snapshot is reset to that of the original disk data (primary) on which the snapshot was based. That is, the access time of the mirror snapshot data continues to match that of the data on the primary disk until the mirror is again synchronized with the primary. This is normal Advanced Client behavior, and applies to local as well as offhost backups.

Simplified View of Snapshot Sequence And Access Time



FlashBackup Configuration

5

This chapter describes FlashBackup and explains how to configure FlashBackup policies.

The following topics are covered in this chapter:

- ◆ [FlashBackup Capabilities](#)
- ◆ [Restrictions](#)
- ◆ [Configuring a FlashBackup Policy](#)
- ◆ [Configuring FlashBackup in the Earlier Manner \(UNIX only\)](#)



FlashBackup Capabilities

FlashBackup is a policy type that combines the speed of raw-partition backups with the ability to restore individual files. The features that distinguish FlashBackup from other raw-partition backups and standard file system backups are these:

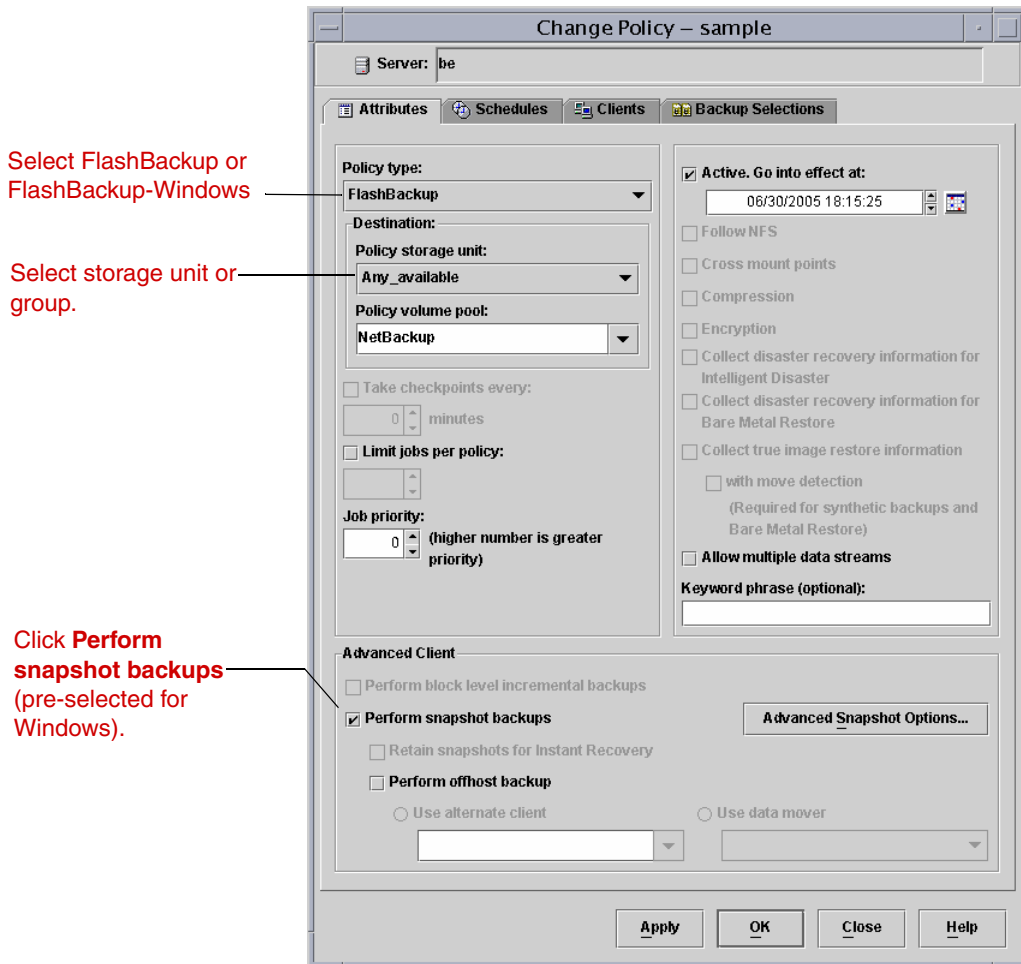
- ◆ Increases backup performance over standard file-ordered backup methods. For example, a FlashBackup of a file system completes in substantially less time than other types of backup, if the file system contains a very large number of files and most of the file system blocks are allocated.
- ◆ Individual files can be restored from raw-partition backups.
- ◆ Backs up the following file systems: VxFS (Solaris and HP), ufs (Solaris), Online JFS (HP), and NTFS (Windows).
- ◆ Supports multiple data streams, to further increase the performance of raw-partition backups when there are multiple devices in the Backup Selections list.

Restrictions

- ◆ FlashBackup policies do not support file systems managed by HSM.
- ◆ FlashBackup does not support VxFS storage checkpoints used by the **VxFS_Checkpoint** snapshot method.
- ◆ FlashBackup supports the following I/O system components: ufs, VxFS, and Windows NTFS file systems, VxVM volumes and LVM volumes, and raw disks. Other components (such as non-VERITAS storage replicators or other non-VERITAS volume managers) are not supported.
- ◆ Note these restrictions for Windows clients:
 - ◆ FlashBackup-Windows policies do not support the backup of Windows system-protected files (the System State, such as the Registry and Active Directory).
 - ◆ FlashBackup-Windows policies do not support the backup of Windows OS partitions that contain the Windows system files (usually C:).
 - ◆ FlashBackup-Windows policies do not support the backup of Windows System database files (such as RSM Database and Terminal Services Database).
 - ◆ FlashBackup-Windows policies do not support include lists (exceptions to client exclude lists).

Configuring a FlashBackup Policy

1. Start the NetBackup Administration Console:
 - ◆ On UNIX, enter: `/usr/openv/netbackup/bin/jnbSA &`
 - ◆ On Windows, click **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**.
2. Click on **Policies**. Double click on the policy (or right-click to create a new one).



3. Select the Policy type: **FlashBackup** for UNIX clients, or **FlashBackup-Windows** for Windows clients.



4. Specify the storage unit.

FlashBackup and FlashBackup-Windows policies support both tape storage units and disk storage units.

5. Select a snapshot method in one of the following ways:

- ◆ Click **Perform snapshot backups** on the Attributes tab.

If this is a new policy, NetBackup will select a snapshot method when the backup starts.

If this is a copy of a policy that was configured for a particular snapshot method, click the **Advanced Snapshot Options** button and set the snapshot method to **auto**. NetBackup will select a snapshot method when the backup starts.

- ◆ Click **Perform snapshot backups**, click the **Advanced Snapshot Options** button, and then select a snapshot method. For more information, refer to “[Selecting the Snapshot Method](#)” on page 86.

6. UNIX only: if you selected nbu_snap or VxFS_Snapshot as the snapshot method, specify a raw partition as cache, in either of these ways:

- ◆ Use the Host Properties node of the Administration Console to specify the default cache device path for snapshots. Click **Host Properties > Clients**, select the client, then **Actions > Properties, UNIX Client > Client Settings**.
- ◆ Use the Advanced Snapshot Options dialog (see “[How to Enter the Cache](#)” on page 145) to specify the cache used for the client when backed up by this policy.

Note The partition to be used for the cache must exist on all clients included in the policy.

7. To shift the backup I/O to an alternate client, or to a NetBackup media server or third-party copy device (for UNIX clients only), select **Perform offhost backup**. See [step 7](#) on page 80 for more instructions on this option.

For FlashBackup, the **Use data mover** option is supported for UNIX clients only.

8. To reduce backup time when more than one raw partition is specified in the Backup Selections list, select **Allow multiple data streams**.

9. Use the Schedules tab to create a schedule.

FlashBackup policies support full and incremental types only. User backup and archive schedule types are not supported.

Note For FlashBackup and FlashBackup-Windows policies, a full backup backs up all blocks in the disk or raw partition selected in the Backup Selections tab (see next step). An incremental backup backs up only the blocks associated with the files that were changed since the last full backup.

- 10.** Use the Backup Selections tab to specify the drive letter or mounted volume (Windows) or the raw disk partition (UNIX) containing the files to back up.

Windows examples:

```
\\.\E:  
\\.\E:\mounted_volume
```

Note The drive must be designated exactly as shown above (E:\ is not correct). Backing up the drive containing the Windows system files (usually the C drive) is not supported.

Solaris examples:

```
/dev/rdisk/c1t0d0s6  
/dev/vx/rdisk/volgrp1/vol1
```

HP examples:

```
/dev/rdisk/c1t0d0  
/dev/vx/rdisk/volgrp1/vol1
```

Note On UNIX: The Backup Selections tab must specify the raw (character) device corresponding to the block device over which the file system is mounted. For example, to back up /usr, mounted on /dev/dsk/c1t0d0s6, enter raw device /dev/rdisk/c1t0d0s6. Note the “r” in /rdisk.

Wildcards (such as /dev/rdisk/c0*) are not permitted. Specifying the actual device file name such as /devices/pci@1f,0/pci@1/scsi@3/sd@1,0:d,raw is not supported.

Note Advanced Client policies do not support the ALL_LOCAL_DRIVES entry in the Backup Selections list.

- 11.** Use the Clients tab to select clients to back up.

Each client in the client list must contain *all* the raw partitions that are specified in the Backup Selections tab. Exclude and include lists cannot be used to circumvent this requirement.



Configuring FlashBackup in the Earlier Manner (UNIX only)

Prior to NetBackup 5.0, FlashBackup was a separate product with two built-in snapshot methods: snapctl driver (nbu_snap) for Solaris clients and a VxFS snap driver for HP clients. The configuration procedure for a pre-5.0 FlashBackup policy was different from that in 5.0 and later, as follows:

- ◆ Unless FlashBackup was combined with ServerFree Agent at NetBackup 4.5, the snapshot methods were preset to the snapctl driver for Solaris clients and VxFS snap driver for HP clients.
- ◆ The raw partition to be used as cache for the snapctl driver and VxFS snap driver had to be specified as a CACHE= entry in the policy's Backup Selections list.
- ◆ Using multiple data streams required adding other directives to the policy's Backup Selections list.

The following sections explain how to configure a FlashBackup policy with a CACHE= entry in the policy's Backup Selections list. This means of configuration is provided for backwards compatibility and will be discontinued in a future release.

Snapshot Selection in Earlier Manner

To set up the policy to use the snapctl driver for Solaris clients or VxFS snap driver for HP, leave **Perform snapshot backups** *unselected* on the policy Attributes tab. NetBackup will use nbu_snap (snapctl driver) for Solaris clients or VxFS_Snapshot for HP.

Cache Partition in Earlier Manner

The snapctl driver and VxFS snap driver are copy-on-write snapshot methods requiring a cache partition. For FlashBackup clients prior to NetBackup 5.0, the cache partition was specified in the policy's Files tab as a CACHE = *raw_partition* entry. (The Files tab is now called the Backup Selections tab.)

This section describes how to specify the CACHE entry.

Note CACHE entries are allowed only when the policy's **Perform snapshot backups** option is unselected. If **Perform snapshot backups** is selected, NetBackup will attempt to back up the CACHE entry and the backup will fail.

- ❖ On the policy's Backup Selections tab, specify at least one cache device by means of the CACHE directive. For example:

CACHE=/dev/rdisk/c0t0d0s1


This is the cache partition for storing any blocks that change in the source data while the backup is in progress. CACHE= must precede the source data entry. See following example.

Backup Selections list with CACHE entry

CACHE entry must come before the source data entry.

All entries must specify the raw device, such as /dev/rdisk/c0t0d0s1.

Do not use the actual file name; must be the link form of cxtxdxsx.



Please note:

- ◆ Specify the *raw* device, such as /dev/rdisk/c1t0d0s6. Do not specify the block device, such as /dev/dsk/c1t0d0s6.
- ◆ Do not specify the actual device file name. The following, for example, is not allowed:
/devices/pci@1f,0/pci@1/scsi@3/sd@1,0:d,raw
- ◆ Wildcards (such as /dev/rdisk/c0*) are not allowed.
- ◆ The CACHE entry must precede the entry for the source data you want to back up.
- ◆ All entries in the Backup Selections list, including the source data to back up, must be the full path name of a *raw* device in the form:

On Solaris: /dev/rdisk/cxtxdxsx

On HP: /dev/rdisk/cxtxdx

where x is an integer.

- ◆ When using multiple data streams, you can include multiple entries in the Backup Selections list.

For example:

CACHE=/dev/rdisk/c1t4d0s0

/dev/rdisk/c1t4d0s7

CACHE=/dev/rdisk/c1t4d0s1



```
/dev/rdsk/c1t4d0s3  
/dev/rdsk/c1t4d0s4
```

Requirements for the cache partition

- ◆ Must reside on the same host as the raw partitions containing the source data that you are backing up.
- ◆ Cannot be the raw partition being backed up.
- ◆ Cannot be a mounted file system. Any data configured on this device may be overwritten by the copy-on-write process.
- ◆ On Solaris, the same cache partition may be used simultaneously by multiple backups (two policies can use the same cache partition at the same time). On HP, multiple backups cannot use the same cache partition simultaneously. If multiple policies list the same cache partition on HP systems, backups naming those policies must run at different times to prevent a failure.
- ◆ The cache partition must have enough space to hold all the writes to the source data that may occur during the backup. Backups during off-peak hours normally require a smaller cache than those during peak activity. See [“Sizing the Cache Partition”](#) on page 143.

Using Multiple Data Streams

For multiple data streams, certain directives must be added to the policy’s Backup Selections tab.

- ◆ The number of backups started depends on the directives in the Backup Selections tab.
- ◆ The number of backups that can run concurrently depends on the number of available drives in the storage units and the maximum jobs parameters (for example, **Limit jobs per policy**).

Note Only one data stream is created for each physical device on the client. You cannot include the same partition more than once in the Backup Selections list.

The directives that you can use in the Backup Selections list for a FlashBackup policy are:

NEW_STREAM

CACHE=*value* (this directive is required; see [“Cache Partition in Earlier Manner”](#) on page 104)

UNSET

UNSET_ALL

Each backup begins as a single stream of data. The start of the Backup Selections list up to the first `NEW_STREAM` directive (if any) is the first stream. Each `NEW_STREAM` entry causes NetBackup to create an additional stream or backup.

Note that all file paths listed between `NEW_STREAM` directives are in the same stream.

The Backup Selections list in the following example generates four backups:

On Solaris systems:

- ❶ `CACHE=/dev/rdisk/c1t3d0s3`
`/dev/rdisk/c1t0d0s6`
- ❷ `NEW_STREAM`
`/dev/rdisk/c1t1d0s1`
- ❸ `NEW_STREAM`
`UNSET CACHE`
`CACHE=/dev/rdisk/c1t3d0s4`
`/dev/rdisk/c1t2d0s5`
`/dev/rdisk/c1t5d0s0`
- ❹ `NEW_STREAM`
`UNSET CACHE`
`CACHE=/dev/rdisk/c0t2d0s3`
`/dev/rdisk/c1t6d0s1`

On HP systems:

- ❶ `CACHE=/dev/cache_group/rvol1c`
`/dev/vol_grp/rvol1`
- ❷ `NEW_STREAM`
`UNSET CACHE`
`CACHE=/dev/cache_group/rvol2c`
`/dev/vol_grp/rvol2`
- ❸ `NEW_STREAM`
`UNSET CACHE`
`CACHE=/dev/cache_group/rvol3c`
`/dev/vol_grp/rvol3`
`/dev/vol_grp/rvol3a`
- ❹ `NEW_STREAM`
`UNSET CACHE`
`CACHE=/dev/cache_group/rvol4c`
`/dev/vol_grp/rvol4`

1. The first stream is generated automatically and a backup is started for `/dev/rdisk/c1t0d0s6` (Solaris) or `/dev/vol_grp/rvol1` (HP). The `CACHE=` entry sets the cache partition to `/dev/rdisk/c1t3d0s3` (Solaris) or `/dev/cache_group/rvol1c` (HP).
2. The first `NEW_STREAM` directive starts a second stream to back up `/dev/rdisk/c1t1d0s1` (Solaris) or `/dev/vol_grp/rvol2` (HP). On Solaris systems, this backup uses the same cache partition. On HP systems, a different cache partition must be defined for each stream (`CACHE=/dev/cache_group/rvol2c`).
3. The second `NEW_STREAM` starts a backup for `/dev/rdisk/c1t2d0s5` and `/dev/rdisk/c1t5d0s0` (Solaris) or `/dev/vol_grp/rvol3` and `/dev/vol_grp/rvol3a` (HP). These two partitions are backed up serially within the stream. In addition, the `UNSET CACHE` directive unsets the previous cache setting and the `CACHE=` directive sets a new cache partition for this backup.
4. The last `NEW_STREAM` directive starts a backup for `/dev/rdisk/c1t6d0s1` (Solaris) or `/dev/vol_grp/rvol4` (HP). Like the third stream, this one also unsets the cache directive and defines a new cache partition.



As shown in the example, policy-specific directives such as `CACHE` are passed to the client with the current stream and all subsequent streams, until the directive is redefined or unset.

If the directive is encountered again, its value is redefined.

An `UNSET` or `UNSET_ALL` directive unsets a directive that was previously defined in the Backup Selections list.

- ◆ `UNSET` unsets a policy-specific directive so it is not passed with additional streams. The directive that was unset can be defined again later in the Backup Selections list to be included in the current or a later stream.
- ◆ `UNSET_ALL` has the same effect as `UNSET` but affects all policy-specific directives that have been defined up to this point in the Backup Selections list. If used, `UNSET_ALL` must appear immediately after the second or later `NEW_STREAM` directive.

NAS Snapshot Configuration

This chapter explains how to configure NetBackup policies to create snapshots of client data on NAS (NDMP) hosts, including backups to SnapVault.

Note NetBackup for NDMP software is required in addition to Advanced Client.

The following topics are covered in this chapter:

- ◆ [NAS Snapshot Overview](#)
- ◆ [SnapVault Overview](#)
- ◆ [Notes and Prerequisites](#)
- ◆ [Setting up a Policy for NAS Snapshots](#)
- ◆ [Configuring SnapVault](#)
- ◆ [Notes on SnapVault](#)
- ◆ [NAS Snapshot Naming Scheme](#)



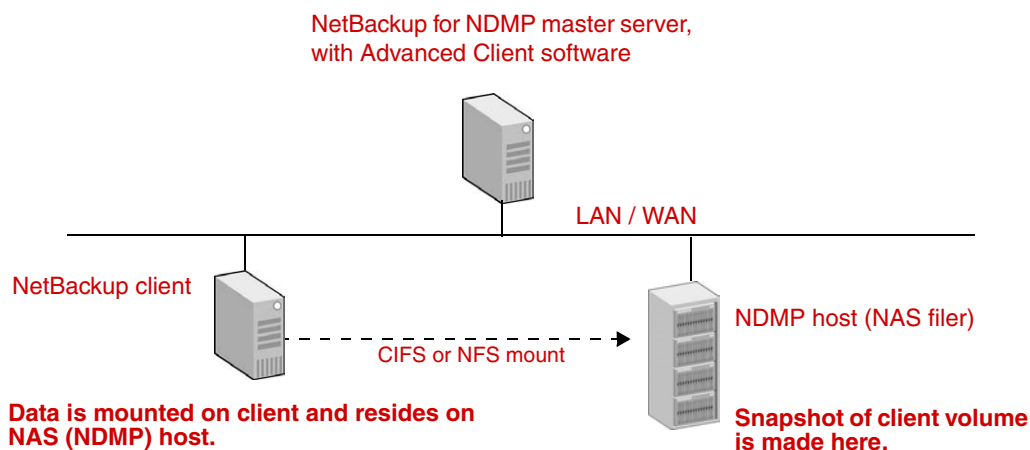
NAS Snapshot Overview

By means of the snapshot feature of Advanced Client and the NDMP V4 snapshot extension, NetBackup can make snapshots of client data on a NAS (NDMP) host. The client data must reside on the NAS host and be mounted on the client by means of NFS on UNIX or CIFS on Windows.

A NAS snapshot is a point-in-time disk image. Snapshots can be retained on disk as long as desired. The data can be efficiently restored from disk by means of the Advanced Client Instant Recovery feature.

See the following diagram for an overview.

NDMP snapshot environment



In NetBackup policy, enter:

For Windows client:

`\\ndmp_hostname\volume`

For UNIX client:

`/NFS_mountpoint`

NOTE: Windows pathnames must use the Universal Naming Convention (UNC).

NetBackup creates snapshots on the NAS-attached disk only, not on storage devices attached to the NetBackup server or the client.

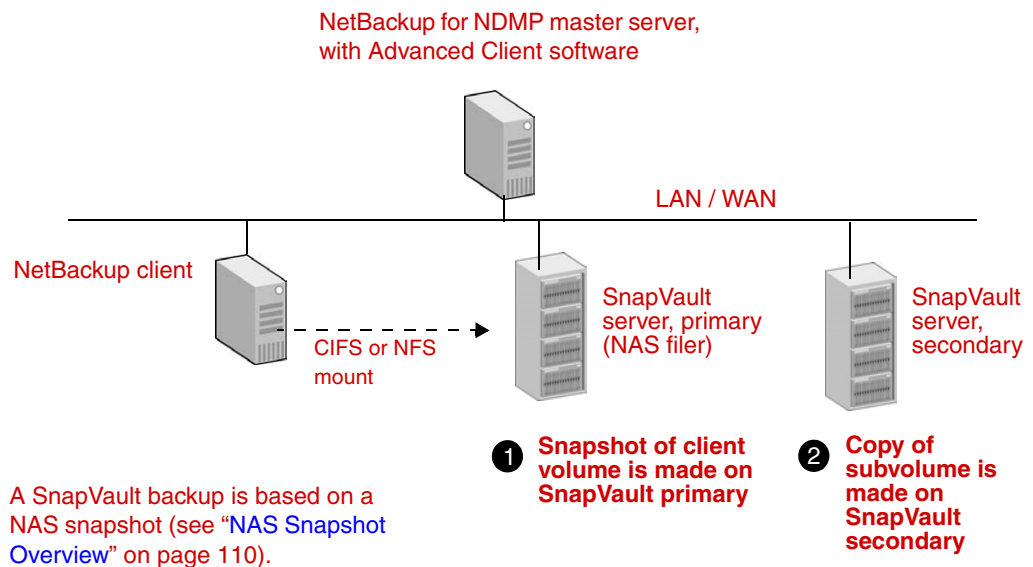
SnapVault Overview

In addition to making a snapshot of client data on the NAS host, NetBackup can also copy the NAS snapshot data to a disk-based SnapVault secondary host for additional security, remote office backup convenience, and speed of restore. In this case, the NAS filer containing the NAS snapshot is the *primary* host, and the SnapVault server containing a disk backup of the snapshot is the *secondary* host. SnapVault backups can be made at frequent intervals and can be retained on disk as long as desired.

SnapVault is an open protocol that NAS vendors can adopt. For a list of NAS vendors that NetBackup currently supports for SnapVault, refer to the *NetBackup Advanced Client Configuration and Compatibility* online pdf (see the preface of this manual for help accessing that document).

See the following diagram for an overview.

SnapVault environment



**Note: on the primary, the snapshot is made of a volume;
On the secondary, the SnapVault is made of a subvolume (qtree).**



Notes and Prerequisites

These notes apply to backups made with the `NAS_Snapshot` method, which is a requirement for SnapVault backups. For additional items unique to SnapVault, see [“Configuring SnapVault”](#) on page 116.

- ◆ Snapshots of NAS host data are supported for NetBackup clients running Windows (32 and 64-bit) and Solaris.
- ◆ Note these software and licensing requirements:
 - ◆ On the NetBackup server: both NetBackup for NDMP and Advanced Client software must be installed and licensed. In addition, a NetBackup for NDMP license must be purchased for each NDMP host.
 - ◆ NetBackup clients that will be used to perform backups must have Advanced Client installed.
 - ◆ On NetBackup clients for Oracle: NetBackup for Oracle database agent software must be installed on all clients.
- ◆ The NAS host must support NDMP protocol version V4 and the NDMP V4 snapshot extension, with additional changes made to the snapshot extension. The *NetBackup Advanced Client Configuration and Compatibility* online pdf contains a list of NAS vendors that NetBackup supports for NAS snapshots, and includes requirements specific to your NAS vendor. Refer to the preface of this manual for help accessing that document.
- ◆ NetBackup must have access to each NAS host on which a NAS snapshot will be created. To create this authorization, you can use either of the following:
 - ◆ In the NetBackup Administration Console: the **Media and Device Management > Devices > NDMP Hosts** option, or the NetBackup Device Configuration wizard.OR
- ◆ The following command:

```
tpconfig -add -nh ndmp_host -user_id user_ID -password password
```

For SnapVault, if you use the `tpconfig` command, you must also enter the following: `tpautoconf -verify -nh ndmp_host`
- ◆ The client data must reside on a NAS host and be mounted on the client by means of NFS on UNIX or CIFS on Windows. For NFS mounts, the data must not be auto-mounted, but must be hard (or manually) mounted.
- ◆ For NAS snapshot backups, and for SnapVault backups, you must create an Advanced Client policy with the `NAS_Snapshot` method as explained under [“Setting up a Policy for NAS Snapshots.”](#)

- ◆ On Windows clients, when restoring files from a backup made with the NAS_Snapshot method, the “NetBackup Client Service” must be logged in as the Administrator account, not as the local system account. The Administrator account allows NetBackup to view the directories on the NDMP host to which the data will be restored. If you attempt to restore files from a NAS_Snapshot (made with or without SnapVault) and the NetBackup Client Service is logged in as the local system account, the backup fails.
 - a. In Windows Services, double-click the NetBackup Client Service.
 - b. Check the Log On tab: if the service is not logged in as Administrator, stop the service.
 - c. Change the log in to the Administrator account and restart the service.
 - d. Retry the restore.

Setting up a Policy for NAS Snapshots

This section explains how to set up a policy for making snapshots of NAS data and for SnapVault. For additional SnapVault requirements, see “[Configuring SnapVault](#)” on page 116.

As an alternative to this procedure, you can use the Snapshot Policy Configuration wizard in the NetBackup Administration Console.

1. Start the NetBackup Administration Console on the NetBackup for NDMP server as follows:

On Windows: from the Windows **Start** menu, select **Programs, VERITAS NetBackup, NetBackup Administration Console**.

On UNIX, enter the following:

```
/usr/opensv/netbackup/bin/jnbSA &
```

2. Click on **Policies** in the left pane.
3. For **Policy type**, select Standard for UNIX clients, MS-Windows-NT for Windows clients, or Oracle for UNIX clients configured in an Oracle database.

Note The NDMP policy type is not supported for snapshots or SnapVault in this 6.0 release.



4. For **Storage unit**, select **Any_available** if this policy is for a NAS snapshot but not for SnapVault. Note the following:
 - ◆ Although the policy cannot execute without a specified storage unit, NetBackup does not use the storage unit for a non-SnapVault snapshot. The snapshot is created on disk regardless of which storage unit you select. Note that the storage unit you select is not reserved, so it is free for use by other policies.
 - ◆ For Oracle policies, the policy uses the storage unit you specify, but only for backing up archive logs and control files.
 - ◆ **For SnapVault:** a SnapVault disk storage unit must be defined. Once defined, this storage unit can be specified here or on the policy **Schedules** tab. See [“Create a Storage Unit for SnapVault”](#) on page 118.
5. Select **Perform snapshot backups** and **Retain snapshots for instant recovery**.
6. Select **Perform offhost backup** and **Use data mover**.
7. From the pull-down under **Use data mover**, pick **Network Attached Storage**.

When the policy executes, NetBackup will automatically select the NAS_Snapshot method for creating the snapshot.

As an alternative, you can manually select the NAS_Snapshot method using the Advanced Snapshot Options dialog from the policy Attributes display.

Note When selecting NAS_Snapshot for SnapVault backups, the **Maximum Snapshots (Instant Recovery only)** parameter determines how many snapshots can be kept on the SnapVault primary, not how many SnapVault copies are kept on the SnapVault secondary. To set the number of SnapVault copies kept on the secondary, use the **Retention** field on the policy Schedule tab (see [step 8](#)).

8. On the **Schedule** Attributes tab, select the following:

Instant Recovery

Choose **Snapshots only**. The other option (**Snapshots and copy snapshots to a storage unit**) does not apply to NAS_Snapshot or SnapVault.

Override policy storage unit

If the correct storage unit was not selected on the Attributes tab, select it here.

For a backup to a SnapVault server, select the disk storage unit created for SnapVault (see [“Create a Storage Unit for SnapVault”](#) on page 118).

Retention

Determines the retention period for SnapVault copies on the SnapVault secondary.

9. For the **Backup Selections** list, specify the directories, volumes, or files *from the client perspective*, not from the NDMP host perspective. For example:
- ◆ On a UNIX client, if the data resides in `/vol/vol1` on the NDMP host `nas1`, and is NFS mounted to `/mnt2/home` on the UNIX client, specify `/mnt2/home` in the Backup Selections list.
 - ◆ On a Windows client, if the data resides in `/vol/vol1` on the NDMP host `nas1`, and is CIFS mounted (mapped to) `\\nas1\vol\vol1` on the Windows client, specify `\\nas1\vol\vol1` in the Backup Selections list.
 - ◆ Windows pathnames must use the Universal Naming Convention (UNC), in the form `\\server_name\share_name`.
 - ◆ The client data must reside on a NAS host and be mounted on the client by means of NFS on UNIX or CIFS on Windows. For NFS mounts, the data must be manually mounted by means of the `mount` command, not auto-mounted.
 - ◆ All paths for a given client in the policy must be valid, or the backup will fail.
 - ◆ The `ALL_LOCAL_DRIVES` entry is not allowed in the Backup Selections list.
 - ◆ **For SnapVault:** you must specify subvolumes (*qtrees* in Network Appliance Data ONTAP), not volumes. For example: if the subvolume is `/vol/vol1/subvol1` on the NDMP host, and is mounted to `/mnt/my_volume/subvol1` on the client, specify `/mnt/my_volume/subvol1` in the Backup Selections list. Each subvolume to be backed up must be entered in the Backup Selections list. For example:

```
/mnt/my_volume/subvol1  
/mnt/my_volume/subvol2  
/mnt/my_volume/subvol3
```

Note For SnapVault: the subvolumes *must be shared* on the NAS filer, by means of NFS on UNIX or CIFS on Windows. In the above example entry (`/mnt/my_volume/subvol1`), `/subvol1` must be shared.

Note also: placing `/mnt/my_volume/subvol1/dir1` in the Backup Selections list will not work, unless `dir1` is also shared.



Configuring SnapVault

This section pertains to SnapVault only.

SnapVault Restrictions

Note For a list of supported NAS systems, and for additional requirements pertaining to your NAS equipment, refer to the *NetBackup Advanced Client Configuration and Compatibility* online pdf (see the preface of this manual for help accessing that document).

- ◆ The Windows Server Appliance Kit (SAK) is not supported.
- ◆ Open Systems SnapVault (OSSV) is not supported.

SnapVault Prerequisites

In addition to completing all requirements for the NAS_Snapshot method, do the following before starting a SnapVault backup.

- ◆ The NetBackup server must have the NAS SnapVault Option license.
- ◆ The SnapVault primary and SnapVault secondary systems must be licensed as primary and secondary by your NAS vendor. A single host cannot have both licenses. Refer to the documentation provided by your NAS vendor.
- ◆ Mount the SnapVault primary subvolumes (qtrees) on the NetBackup client. These are the subvolumes that you want to back up to a SnapVault secondary. The subvolumes must be shared by means of NFS on UNIX or CIFS on Windows.
- ◆ Create a SnapVault storage unit.
- ◆ Configure a NetBackup policy for the NAS_Snapshot method (explained in previous section).
- ◆ Enable NetBackup access to the SnapVault primary system.
- ◆ Enable access between the SnapVault primary and SnapVault secondary systems.

Mount Primary Subvolume On Client

NetBackup's SnapVault feature supports subvolume-to-subvolume transfers only. That is, the smallest unit of data that can be selected for SnapVault backup is a subvolume. (In Network Appliance Data ONTAP, a subvolume is a *qtree*.) Prior to running a SnapVault backup, you must make sure that the subvolume(s) on the SnapVault primary are mounted on the NetBackup client and shared.

- ◆ For the subvolume on the SnapVault primary, mount it on the NetBackup client as follows:

For a UNIX client, create an NFS (Network File System) mount. For example:

`/NFS_mountpoint`

For a Windows client, create a CIFS (Common Internet File System) mount. For example:

`\\ndmp_hostname\share_name`

Note Windows pathnames must use the Universal Naming Convention (UNC).

See [step 9](#) on page 115 for instructions on specifying the mounted volume in the NetBackup policy.

Enable NetBackup Access to the SnapVault Primary

To make the SnapVault, NetBackup must have access to the SnapVault primary system.

1. On the NetBackup server: under **Media and Device Management > Devices**, click on **NDMP Hosts**. Under **Actions**, choose **New > NDMP Host**.
2. Enter the name of the SnapVault primary (NDMP host) and use the NDMP Host dialog to enter the credentials (username and password) that the NetBackup server will use to access the SnapVault primary.

For more detail on this procedure, refer to “Authorizing Access to the NDMP Host” in the *NetBackup for NDMP System Administrator’s Guide*.

Enable Access Between SnapVault Primary and SnapVault Secondary

The SnapVault primary machine and SnapVault secondary must be able to access each other. The procedure for creating this access is specific to each NAS vendor. For assistance, refer to the documentation supplied by your NAS vendor. Additional information is also available in the *NetBackup Advanced Client Configuration and Compatibility* online pdf (see the preface of this manual for help accessing that document).



Create a Storage Unit for SnapVault

Backups to a SnapVault server require a special disk storage unit, to be defined as follows.

1. In the NetBackup Administration Console, go to **Storage Units**, and from the **Actions** menu select **New > Storage Unit**. The New Storage Unit dialog appears.

The screenshot shows the 'New Storage Unit' dialog box with the following fields and annotations:

- Storage unit name:** SnapVault1 (Annotation: Specify storage unit name.)
- Storage unit type:** Disk (Annotation: Select **Disk** as storage unit type.)
- On demand only:** ☒ (Annotation: Select **SnapVault** as the disk type.)
- Disk type:** SnapVault (Annotation: Select **SnapVault** as the disk type.)
- Media server:** media_server1 (Annotation: Select the media server for the SnapVault backup.)
- SnapVault server:** SnapVault_secondary (Annotation: Select the SnapVault secondary system on which to make the backup.)
- Absolute pathname to volume:** /vol/QE2 (Annotation: Select the volume in which to store the SnapVault.)
- Properties and Server Selection** (Section header)
- Maximum concurrent jobs:** 1 (Annotation: Select Transfer throttle, if desired.)
- Reduce fragment size to:** 524288 Megabytes
- High water mark:** 98 %
- Transfer throttle:** 9999999 Kilobytes/second
- Buttons:** OK, Cancel, Help

2. Specify the following:

Storage unit name

A unique name for this storage unit.

Storage unit type

Select **Disk**.

Note The 6.0 release of NetBackup does not support tape backups of SnapVault.

On demand only

This is required. **On demand only** prevents NAS_Snapshot policies that are configured with the Any_available storage unit from using this SnapVault storage unit. If **On demand only** is not selected, NAS_Snapshots not intended for SnapVault may end up being transferred to the SnapVault secondary.

Disk type

Select **SnapVault**.

Note Do not select NearStore. The NearStore storage unit type does not support SnapVault backups.

Media server

Select the NetBackup media server that will control the SnapVault storage unit.

SnapVault server

Select the SnapVault secondary system on which to create the SnapVault.

Absolute pathname to volume

From the pull-down, select the volume on the SnapVault secondary in which to create the SnapVault.

Properties

For the SnapVault secondary volume, click **Properties** for a listing of its total capacity and available space.

Transfer throttle

This option makes it possible to limit the amount of network bandwidth (in kilobytes per second) used for the SnapVault transfer, in case bandwidth needs to be reserved for other applications. Zero (the default) means *no network bandwidth limit* for the SnapVault transfer: SnapVault will use all available bandwidth.

A value greater than zero indicates a transfer speed for SnapVault in kilobytes per second. A value of 1, for instance, would set a transfer speed limit for SnapVault of 1 kilobyte per second, which is a *very* slow transfer rate.

3. Click **OK**.



Relationships Between Primary and Secondary SubVolumes

Note the following about SnapVault primary and secondary subvolumes (qtrees):

When does NetBackup perform a full SnapVault?

Changed blocks only

If an initial SnapVault has already occurred between the primary subvolume and the corresponding subvolume on the SnapVault secondary system, all subsequent SnapVaults from the primary to the secondary subvolume will consist of changed blocks only.

Full SnapVault

If an initial SnapVault has not yet occurred, a full SnapVault is performed: the entire contents of the primary subvolume are copied to the SnapVault on the secondary. The time required to make the SnapVault depends on the size of the primary subvolume and the available network bandwidth.

When does NetBackup create a new subvolume for the SnapVault?

If the subvolume on the SnapVault primary system has a relationship with *only one subvolume* in the SnapVault secondary's volume, then NetBackup uses the existing subvolume for the SnapVault. No new subvolume is created on the SnapVault secondary system. Since the primary-secondary relationship already exists and a full SnapVault has already occurred, NetBackup creates the SnapVault from changed blocks only.

In all other cases, NetBackup creates a new subvolume on the SnapVault secondary. For example, a new subvolume is created if:

- ◆ There are no subvolumes in the secondary volume
- ◆ A relationship exists between the primary subvolume named in the policy and *two or more* subvolumes in the secondary volume.

In both of these cases, NetBackup creates a new subvolume for the SnapVault. The SnapVault secondary is a complete copy of the SnapVault primary's subvolume.

Notes on SnapVault

- ◆ You can restore SnapVault backups with a normal NetBackup restore or with Point in Time Rollback. See the NetBackup Backup, Archive, and Restore online help for assistance with normal restores. For Point in Time Rollback, see [“Instant Recovery: Snapshot Rollback”](#) on page 199, in addition to the Point in Time Rollback notes below.
- ◆ If you back up a subvolume (qtree) at the root level (such as `nas_host1:vol/vol0/sub_vol1`), the backup will succeed but the restore will fail. You cannot restore to the NAS filer’s root volume. You must restore to an alternate location.

Note the following about Point in Time Rollback:

- ◆ For Point in Time Rollback, the entire subvolume (qtree) is restored, not individual files.
- ◆ For Point in Time Rollback, a SnapVault restore *always goes to a new subvolume on the primary host*, never overwriting an existing subvolume. As a result:
 - ◆ In the NetBackup **Restore Marked Files** dialog, verification is always skipped, regardless of whether or not **Skip verification and force rollback** is selected.
 - ◆ In the NetBackup **Restore Marked Files** dialog, the settings **Restore everything to its original location** and **overwrite existing files** will be ignored.

The format of the new subvolume name is the following:

`mountpointname_restore.time_stamp`

For example: `subvol1_restore.2005.05.19.10h49m04s`.

To use the restored, new subvolume, you could unmount and rename the original subvolume (the subvolume that was not overwritten), rename the new subvolume with the name of the original, and mount the new subvolume on the client.



NAS Snapshot Naming Scheme

Below is the format of a NAS snapshot name:

`NAS+NBU+PFI+client_name+policy_name+<sr|sv>+volume_name+date_time_string`

Note:

- ◆ The snapshot name always begins with NAS+NBU+PFI+
- ◆ The plus sign (+) separates name components.
- ◆ NAS snapshots reside on the NDMP host (NAS filer).

For example:

`NAS+NBU+PFI+sponge+snapvault+sv+Vol_15G+2005.05.31.13h41m41s`

Where:

Client name = sponge

Policy name = snapvault

sr|sv = indicates whether the snapshot was created for a NAS snapshot (sr) or for SnapVault (sv). This allows for separate snapshot rotations on a policy basis, for NAS snapshots and SnapVault snapshots.

Volume name = Vol_15G

Date/Time = 2005.05.31.13h41m41s

Instant Recovery Configuration

This chapter explains how to configure a policy that uses the Instant Recovery feature.

The following topics are covered in this chapter:

- ◆ [Instant Recovery Capabilities](#)
- ◆ [Requirements](#)
- ◆ [Restrictions](#)
- ◆ [Instant Recovery Overview](#)
- ◆ [Configuring a Policy for Instant Recovery](#)
- ◆ [Configuring VxVM](#)
- ◆ [Instant Recovery for Databases](#)



Instant Recovery Capabilities

The Instant Recovery feature of Advanced Client enables high-speed data retrieval from disk by means of the standard NetBackup user interface. Note the following features:

- ◆ Supports NetBackup clients running on Solaris, HP, AIX, Linux, and Windows. The master server can be running on any VERITAS-supported operating system. AIX and Linux client data must reside in a VxFS file system.
- ◆ Uses snapshot technologies to create disk images.
- ◆ Can create a snapshot and a backup to tape, from one policy.
- ◆ Enables random-access (non-sequential) restores of dispersed files from full backups.
- ◆ Enables block-level restore and file promotion from VxFS_Checkpoint snapshots (UNIX). Also enables fast file resync from vxvm and FlashSnap snapshots on Windows.
- ◆ Enables rollback from backups created using VERITAS VxFS Storage Checkpoints, vxvm, FlashSnap, or NAS_Snapshot method.
- ◆ Can restore to a different path or host.
- ◆ Provides resource management by means of a rotation schedule.
- ◆ Supports Oracle, Microsoft Exchange, DB2, and SQL-Server databases.

Requirements

- ◆ For snapshots using Storage Checkpoints, by means of NetBackup's VxFS_Checkpoint method, all Solaris clients must have VxFS 3.4 or later (HP clients require VxFS 3.5 or later, Linux and AIX clients VxFS 4.0 or later) with the Storage Checkpoint feature.
- ◆ For VxVM snapshot volumes on UNIX, clients must have VxVM 3.2 or later with the FastResync feature (Linux and AIX clients require VxVM 4.0 or later). Windows clients must have Storage Foundations for Windows version 3.1.
- ◆ For Instant Recovery with DB2, Oracle, Exchange, or SQL-Server databases, refer to the appropriate NetBackup database agent guide.
- ◆ For replication hosts (using NetBackup's VVR method), clients must have VxVM 3.2 or later with the VERITAS Volume Replicator feature (Linux and AIX clients require VxVM 4.0 or later).

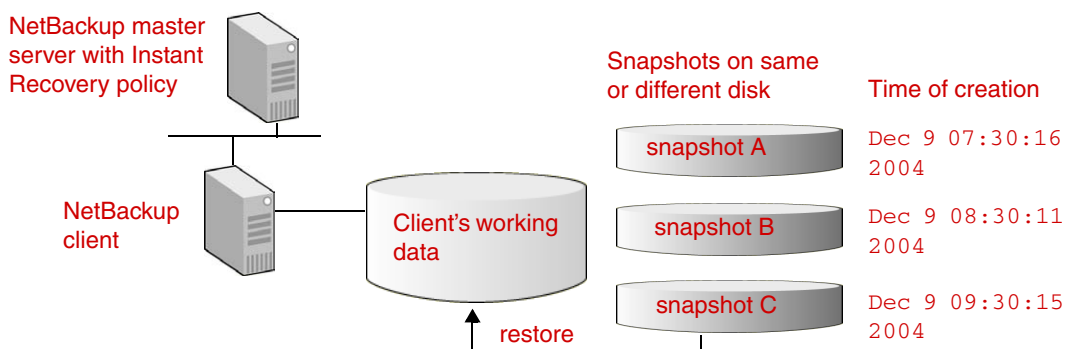
Restrictions

- ◆ For snapshots using Storage Checkpoints, Instant Recovery supports file systems with the Version 4 disk layout or later. Older disk layouts must be upgraded to Version 4 or later.
- ◆ No-data Storage Checkpoints (those containing file system metadata only) are not supported.
- ◆ Instant Recovery snapshots must not be manually removed or renamed, otherwise the data cannot be restored.
- ◆ Instant Recovery does not support the vxvm, FlashSnap, and VVR snapshot methods when used with VxVM volume sets.
- ◆ Alternate client backup is supported in the split-mirror configuration only, using a mirror-type snapshot method (FlashSnap or VVR).
- ◆ For Instant Recovery backups of data configured on VxVM volumes on Windows, the VxVM volume names must be 12 characters or less. Otherwise, the backup will fail.

Instant Recovery Overview

Standard NetBackup can use disks for backup and restore. The Instant Recovery feature of Advanced Client extends that capability by exploiting the speed of snapshots. A snapshot is created with minimal impact on clients' access to data or on the speed of their transactions. If snapshots are made frequently, a user who accidentally deletes or modifies a file can restore the file from the latest snapshot in a matter of seconds.

Instant Recovery from Snapshot on Disk



The on-disk snapshots become point-in-time versions of file systems or volumes, to be kept or discarded as needed. The data can be restored from local disk; no need to mount a tape or other remote storage.

In the figure above, NetBackup Instant Recovery creates snapshot A of the client data on disk. One hour later, as scheduled, NetBackup creates snapshot B, also on disk, followed one hour later when it creates snapshot C. When needed, a user can restore data directly from disk, from the appropriate snapshot.

Note NetBackup Instant Recovery retains the snapshot. The snapshot can be used for restore even if the client has been rebooted.

The following sections provide background information on Storage Checkpoints and snapshot volumes, and describe how NetBackup creates and maintains each type of snapshot.

Snapshot and Backup

An Instant Recovery backup creates a snapshot on disk and optionally backs up the client's data to a storage device. The snapshot is made in one of three places:

- ◆ On the same disk file system that contains the client's original data (requires VxFS file systems, using the VxFS_Checkpoint snapshot method)
- ◆ On a local mirror disk (for VxVM volumes, using the vxvm or FlashSnap snapshot method)
- ◆ On a mirror disk on a replication host (for VERITAS Volume Replication, using the VVR snapshot method)

Maintaining the NetBackup Catalogs

For Instant Recovery backups, NetBackup automatically updates the catalog to keep it correlated with the snapshots on the client. If not kept up-to-date, the catalog might eventually include references to snapshots that no longer exist on the client, due to user activity (snapshot replacement or deletion).

NetBackup includes a maintenance program (`bppficorr`) that can be run manually to update the catalog, if an Instant Recovery snapshot is accidentally deleted or renamed.

For man page information on the `bppficorr` command, refer to the *NetBackup Commands guide*.

VxFS Storage Checkpoints

Storage Checkpoints are a feature of the VERITAS File System (VxFS). They are produced by a copy-on-write mechanism that creates a snapshot by identifying only the file system blocks that have changed since the last checkpoint was taken. When restoring from a storage checkpoint, NetBackup retrieves only the data that changed since the last backup, not the entire file system, if block-level restore is applicable.

For more detail on Storage Checkpoints, refer to the *VERITAS File System Administrator's Guide*. For an introduction to the copy-on-write process, refer to [“How Copy-on-Write Works”](#) on page 232.

A Storage Checkpoint has the following features:

- ◆ Persists after a system reboot or failure.
- ◆ Identifies the blocks that have changed since the last Storage Checkpoint.
- ◆ Shares the same pool of free space as the file system. The number of checkpoints is limited only by the available disk space.
- ◆ Supports mounting a VxFS 4.0 file system over a VxVM 4.0 volume set (multi-device system).

VxVM Split Mirrors

The second type of snapshot that Instant Recovery can make is a *split mirror*. This is a complete disk copy of the client's primary volume. It is stored on separate disk(s), physically independent of the primary volume. Since mirroring requires a complete copy of the primary disk on a separate disk (equal in size to the original disk), it consumes more disk space than a copy-on-write type of snapshot such as a Storage Checkpoint (VxFS_Checkpoint).

For Instant Recovery, the vxvm snapshot method supports split-mirror snapshots.


Snapshot Rotation

Changes made to the files on the client's primary file system are reflected in the Instant Recovery checkpoint or mirror, until the backup policy is again executed. At that point, further changes are reflected in the second checkpoint or pre-allocated mirror, not the first. The first snapshot is retained, in case a restore is needed from that data. When a third backup is executed, the mirrors are “rotated” again and further changes are reflected in the third mirror (if a third mirror was allocated). If a third mirror was not allocated, the changes are reflected in the first mirror, and the data that had been retained on that mirror is overwritten. Note that rotation can be specified in the policy by means of the Maximum Snapshots setting; see [“Maximum Snapshots Setting \(Advanced Snapshot Options dialog\)”](#) on page 128.



The next figure shows rotation of three mirror volumes for Instant Recovery. Rotation occurs each time the Instant Recovery policy is executed.

Mirror Rotation, For Multiple Undo Levels



In the above example, the next Instant Recovery backup will overwrite the mirror snapshot that was made at 12:00 noon.

Configuring Snapshot Deletion

Storage checkpoints consume disk space whenever the primary data changes. Snapshots using mirrors consume another mirror for each snapshot. To avoid running out of disk space or mirrors as snapshots accumulate, some snapshots have to be deleted to make room for new ones.

There are two ways of controlling when Instant Recovery snapshots are deleted: the maximum snapshots setting and the backup retention period.

Maximum Snapshots Setting (Advanced Snapshot Options dialog)

This sets the maximum number of instant recovery snapshots to be retained at one time. When the maximum is reached, the next snapshot causes the oldest to be deleted. If the number of mirrors configured for the snapshots is less than the Maximum Snapshots value, the number of mirrors determines how many snapshots are kept (if the snapshot method uses mirrors).

For deleting snapshots, however, it may be best to use a retention period rather than the maximum snapshots setting (explained under [“Backup Retention Period”](#) on page 129).

Caution If you specify a number that is smaller than the existing number of snapshots, NetBackup deletes the older snapshots until the number of snapshots equals that specified for **Maximum Snapshots**.

Backup Retention Period

This determines when the backup image (snapshot) expires. The shorter the retention period, the sooner the snapshot expires. A short retention period can be useful if your Instant Recovery policy makes snapshots at frequent intervals and the snapshots require a lot of space. Timely expiration of older snapshots will save disk space.

Caution If the retention period is *shorter* than the time between snapshots, each snapshot will expire before the next snapshot occurs. This means that no Instant Recovery backup is available from the moment of expiration until the next snapshot occurs. Avoid this scenario. See [“Examples for Retention Period and Maximum Snapshots”](#) on page 129.

It is usually best to have snapshots deleted when they expire rather than when they exceed the maximum snapshots value. The Maximum Snapshots value can be used as a fallback setting for retention periods of infinity, such as Instant Recovery to tape. It can also serve as a safeguard in case retention periods are not set properly.

The best setting for retention period and maximum snapshots depends many factors. For example:

- ◆ The amount of file system space or number of disk mirrors available for Instant Recovery snapshots
- ◆ How often the primary file system changes
- ◆ How much new data is being added
- ◆ How often Instant Recovery snapshots are created (how often the policy is executed)

Examples for Retention Period and Maximum Snapshots

Example 1 (Using Retention Period)

You have seven VxVM mirrors and want daily FlashSnap Instant Recovery backups, keeping each backup for a week. Configure the backup frequency, retention period, and maximum snapshots value as follows:

Backup frequency = 1 day (set on the policy's **Schedule > Attributes** tab)

Retention period = 1 week (set on the policy's **Schedule > Attributes** tab)

Maximum Snapshots value = 7 (set on the policy's **Advanced Snapshot Options** dialog)



During the first week, a mirror is used each day to take a snapshot. On Monday of the second week, the previous Monday's backup expires (Retention period = 1 week) and its snapshot is deleted to make a mirror available for the Monday backup. The same happens on each of the remaining days of the week and in future weeks. In this case, the Maximum Snapshots value is just a safe guard and is not used.

Example 2 (Using Maximum Snapshots value)

You have three VxVM mirrors and want daily FlashSnap Instant Recovery backups, keeping each backup for a week. Since only three mirrors are available, you cannot keep more than three backups at a time. Configure the backup frequency, retention period, and maximum snapshots value as follows:

Backup frequency = 1 day

Retention period = 1 week

Maximum Snapshots value = 3

On Monday, Tuesday, and Wednesday, a mirror is used each day to take a snapshot. On Thursday, there is no mirror available and no backups have expired (Retention period = 1 week). Because the number of retained backups (3) equals the Maximum Snapshots value of 3, Monday's backup is deleted to release a mirror for the Thursday snapshot. The same happens on each of the following days. In this configuration, the Maximum Snapshots value forces the oldest backup to be deleted.

Retention period must not be less than the period between backups

Improper use of backup frequency and retention period can lead data loss. For example:

Backup frequency = 2 weeks

Retention period = 1 week

In this example, the Instant Recovery backup is made every two weeks, but the backup expires after one week. As a result, during the second week after the backup was made, there are no backups available for Instant Recovery. Unlike tape backup, an Instant Recovery backup that expires is gone.

Configuring a Policy for Instant Recovery

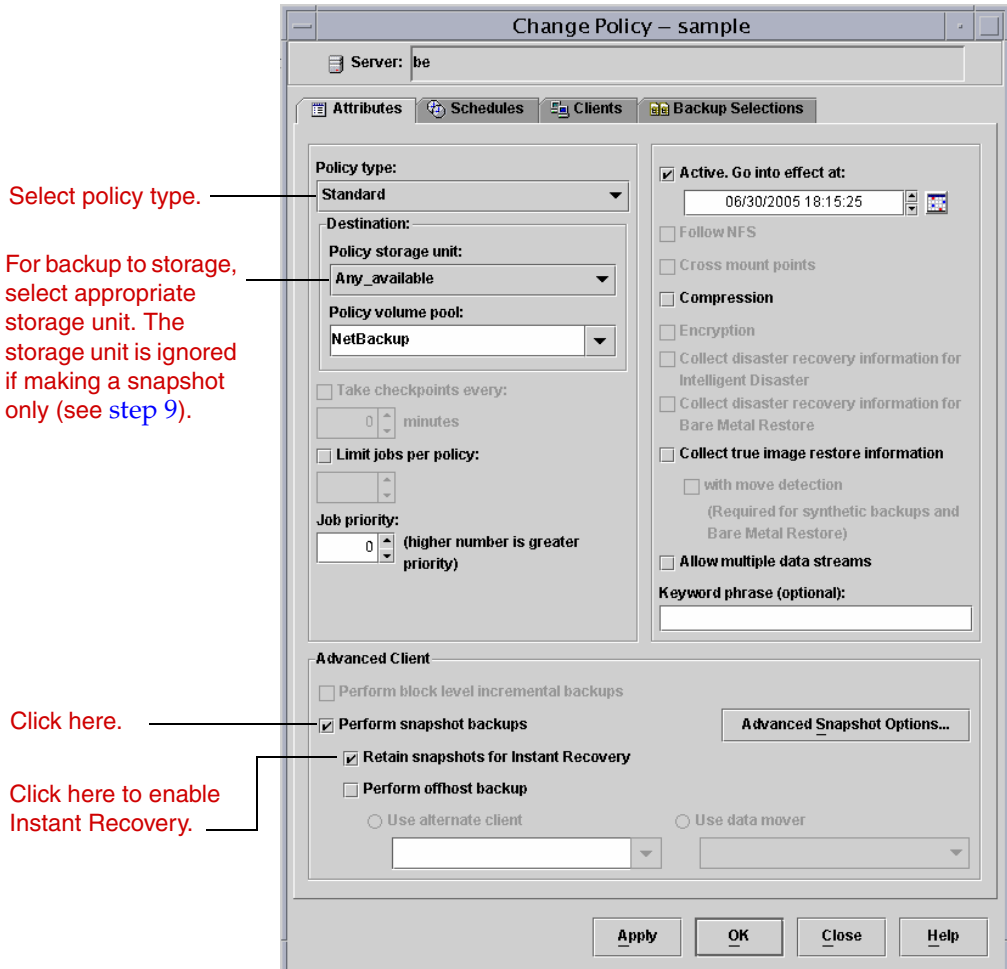
This section explains how to configure a policy for backups enabled for Instant Recovery. As an alternative to this procedure, you can use the **Snapshot Policy Configuration wizard**.

1. Start the NetBackup Administration Console as follows:

On UNIX, enter: **/usr/opensv/netbackup/bin/jnbSA &**

On Windows, click **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**.

2. Click **Policies**. In the **All Policies** pane, open a policy or create a new one.



3. For the policy type, select **Standard, MS-Windows-NT**, or the database agent type appropriate for the client(s).
4. Select a storage unit (disk or tape).

If you select **Snapshots only** on the **Schedules** tab (see [step 9](#)), the storage unit is not used; NetBackup creates a snapshot only.

5. Select **Perform snapshot backups**.
6. Select **Retain snapshots for instant recovery**.

NetBackup retains the snapshot so that Instant Recovery can be performed from the snapshot. A normal backup to storage is also performed, if you do not select **Snapshots only** on the **Schedules** tab ([step 9](#)).

7. To save the settings, click **Apply**.
8. Click the **Advanced Snapshot Options** button to select the snapshot method.

If this is a new policy, you can skip this step to let NetBackup select the method (**auto** is the default).

Select snapshot method.
Choose **auto** if you want
NetBackup to select the
method.

Change values if needed.
If **auto** is the method, only
one snapshot can be
retained at a time.

Snapshot Options – Policy sample

The following selections are optional. If no snapshot method is selected, NetBackup will select a snapshot method at the time of backup.

Snapshot method for this policy:

vxvm

The following parameters can be set for backups performed in this policy that use this snapshot method:

Parameter	Value
Resynchronize mirror in background (not Instant Recovery)	No
Maximum Snapshots (Instant Recovery only)	1

OK Cancel Help

- a. Select a snapshot method from the pull-down list. For creating an Instant Recovery snapshot, the available methods are:
 - ◆ **auto** (UNIX or Windows): NetBackup selects the snapshot method.

- ◆ **FlashSnap** (UNIX or Windows): uses the VxVM FlashSnap feature and VxVM FastResync to create the snapshot. VxVM mirrors must be configured as explained under “[Configuring VxVM](#)” on page 135. Can also use instant snapshots (see “[VxVM Instant Snapshots](#)” on page 151) and, for Windows clients, the Fast File Resync feature of Storage Foundation for Windows 4.1.
- ◆ **NAS_Snapshot** (UNIX or Windows): uses the NDMP V4 snapshot extension to create the snapshot on the NAS-attached disk. Refer to the “[NAS Snapshot Configuration](#)” chapter for help in setting up a policy for NAS_Snapshot.
- ◆ **VxFS_Checkpoint** (UNIX only): uses VxFS Storage Checkpoint to create the snapshot.

A new Storage Checkpoint is created whenever a backup using the VxFS_Checkpoint method is executed.

- ◆ **vxvm** (UNIX or Windows): uses VxVM FastResync to create the snapshot. VxVM mirrors must be configured as explained under “[Configuring VxVM](#)” on page 135. Can also use instant snapshots (see “[VxVM Instant Snapshots](#)” on page 151) and, for Windows clients, the Fast File Resync feature of Storage Foundation for Windows 4.1.
 - ◆ **VVR** (UNIX): creates a snapshot of a VxVM volume on a VERITAS Volume Replication host.
- b. Change parameter values for the method, if needed. The parameters are described under [step 6](#) on page 89.
 - c. When finished, click **OK**.



9. Use the **Schedule** tab to configure a schedule.

Click here for a snapshot only.
 Select **Snapshots and copy snapshots to a storage unit** if you want a normal backup to tape *in addition to* a snapshot for Instant Recovery.
 Note: You must deselect **Snapshots only** if you want to deselect **Retain snapshots for instant recovery** on the policy Attribute tab.

a. For a snapshot only, select **Snapshots only**.

- ◆ If **Snapshots only** is selected, the snapshot is not backed up to tape or other storage. NetBackup creates a snapshot on disk only. This option is required for the **NAS_Snapshot** method.

Note If the snapshot uses the **VxFS_Checkpoint** method or is a **vxvm** space-optimized snapshot, the snapshot is created on the same device as the one containing the original data. In that case, you may want to create another policy to back up the data to a separate device.

- ◆ If **Snapshots and copy snapshots to a storage unit** is selected, NetBackup creates (and retains) a snapshot and backs up the client's data to the storage unit specified for the policy.

b. You can select the retention period for snapshots under **Retention**.

c. Make other schedule selections as desired, and click **OK**.

10. Use the **Backup Selections** tab to enter the files and folders to be backed up.

- ◆ If backing up Oracle database clients, refer to the *NetBackup for Oracle System Administrator's Guide* for instructions.
- ◆ Advanced Client policies do not support the ALL_LOCAL_DRIVES entry in the policy's Backup Selections list.
- ◆ If you use the Backup Policy Configuration wizard, see "[Backup Policy Configuration Wizard](#)" on page 97.

11. Use the **Clients** tab to specify clients to be backed up by this policy.

Configuring VxVM

Note For Instant Recovery backups of data configured on VxVM volumes on Windows, the VxVM volume names must be 12 characters or fewer. Otherwise, the backup will fail.

Before using an Instant Recovery policy for backing up VxVM volumes, one or more mirrors must be created. The primary volumes must be enabled for FastResync. Note that on Windows, FastResync is enabled by default.

Creating a Snapshot Mirror

You can create a snapshot mirror using VxVM commands. As an alternative, you can use the Snapshot Policy Configuration wizard in the NetBackup Administration Console.

Windows

There are two ways to create the snapshot mirror:

- ◆ For a volume that is associated with a drive letter:

```
vxassist snapstart X:
```

where X is the drive letter. This creates a snapshot mirror of the designated drive.

- ◆ For a volume that is not associated with a drive letter:

```
vxdg -g disk_group dginfo
```

This shows information for the specified disk group, including the names of the volumes configured for that group. Create the snapshot by entering the following:

```
vxassist snapstart \Device\HarddiskDmVolumes\disk_group\Volume_name
```

This creates a snapshot mirror of the designated Windows volume.



UNIX

1. Add dco (data change object) logs to the primary volume:

```
/usr/sbin/vxassist -g disk_group addlog volume_name logtype=dco
```

2. Enable FastResync on the volume:

```
/usr/sbin/vxvol -g disk_group set fmr=on volume_name
```

3. Prepare a new mirror for the Instant Recovery snapshot:

- a. Create the mirror:

```
/usr/sbin/vxassist -g disk_group snapstart primary_volume
```

Wait until the mirror is synchronized (status SNAPDONE, or **State** field reads **Ready** in the volume's properties display).

- b. To verify that the mirror is synchronized, enter:

```
/usr/sbin/vxprint -g disk_group -q -t -e 'assoc="primary_volume"'
```

Creating an Instant Snapshot

The following procedure is required for Instant Recovery backups when using the full-sized or space-optimized instant snapshot options in VxVM 4.0. For Instant Recovery backups, VxVM 4.0 or later instant snapshots are supported by Advanced Client's FlashSnap, vxvm and VVR methods.

Creating space-optimized snapshots

A cache object called NBU_CACHE must be created in the disk group containing the volume to be backed up. NetBackup will recognize the cache object and use it to create a space-optimized snapshot.

Note These steps can be performed by the Snapshot Policy Configuration wizard when creating a new policy.

1. Create the parent volume:

```
/usr/sbin/vxassist -g disk_group make volume size layout=layout  
logtype=dco dcoversion=20 [drl=no|sequential|yes] [ndcomirror=number]  
[fastresync=on]
```

Where:

- ◆ Brackets [] indicate optional items.
- ◆ `make volume` specifies the name of the volume snapshot.

2. Create the cache object:

```
/usr/sbin/vxassist -g disk_group make cache_volume size  
layout=layout init=active
```

3. Label the cache object:

```
/usr/sbin/vxmake -g disk_group cache NBU_CACHE  
cachevolname=cache_volume
```

4. Enable the cache object:

```
/usr/sbin/vxcache -g disk_group start NBU_CACHE
```

Creating full-sized snapshots

Unlike the space-optimized snapshot, VxVM full-sized instant snapshots cannot be created by NetBackup: you must create them prior to running the backup, as explained in the following procedure. You must create one full-sized instant snapshot for each backup you are going to run.

Note These steps can be performed by the Snapshot Policy Configuration wizard when creating a new policy.

1. Create the parent volume:

```
/usr/sbin/vxassist -g disk_group make volume length length layout=layout  
logtype=dco dconversion=20 [drl=no|sequential|yes] [ndcomirror=number]  
[fastresync=on]
```

Where:

- ◆ Brackets [] indicate optional items.
- ◆ `make volume` specifies the name of the volume snapshot.

2. Create a volume for a full-sized instant snapshot:

a. Determine the required size for the snapshot volume:

```
# LEN='vxprint -g disk_group -F%len volume'
```

b. Find the name of the DCO volume:

```
# DCOVOL='vxprint -g disk_group -F%dconame volume'
```



- c. Discover the DCO volume's region size (in blocks):

```
# RSZ='vxprint -g disk_group -F%regionsz $DCOVOL'
```

- d. Create a volume called *volumename_NBU*, of the required size and redundancy.

Note The volume name must end with *_NBU*. In the following example, the volume is named *SNAP_vol1_NBU*.

```
vxassist -g disk_group make SNAP_vol1_NBU $LEN layout=mirror
nmirror=number logtype=dc0 drl=no dconversion=20
ndcomirror=number regionsz=$RSZ init=none
[storage attributes ...]
```

The number for *nmirror* should equal the number for *ndcomirror*.

- e. Create the mirror:

```
vxsnap -g disk_group make source=volume/snapvol=SNAP_vol1_NBU/syncing=on
```

3. Set the **Maximum Snapshots (Instant Recovery only)** value on the NetBackup Advanced Snapshot Options dialog.

Using the VxVM Graphical User Interface

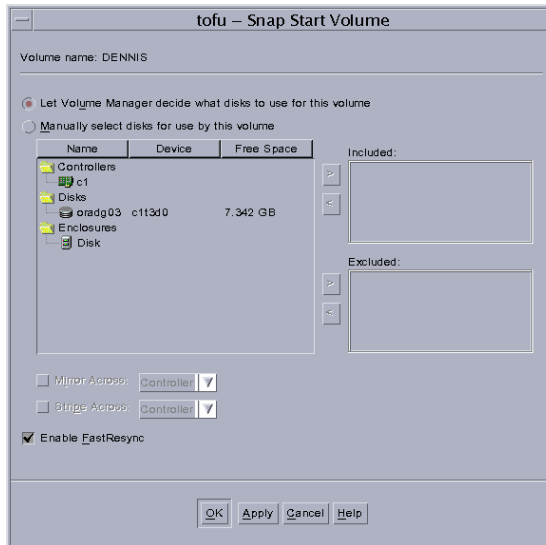
These are the steps for using the VxVM graphical user interface to configure VxVM mirrors for Instant Recovery backups.

Using VERITAS Enterprise Administrator (VxVM 3.5)

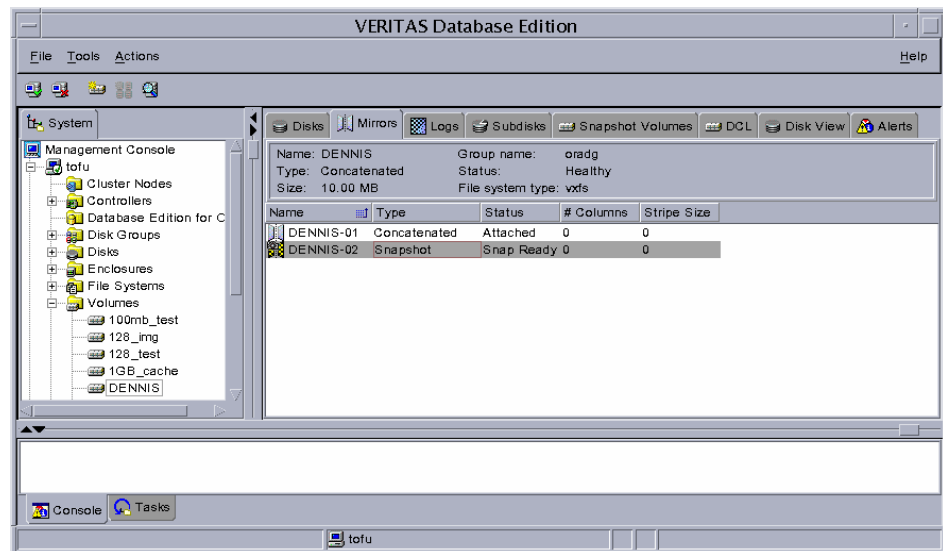
▼ To prepare a new mirror

1. Make sure that FastResync is enabled on the primary VxVM volume.
 - a. In the VEA console, right-click on the volume and click **Properties** from the pop-up menu.
The **FastResync** field states whether or not FastResync is enabled.
 - b. Click **Cancel**.
 - c. If FastResync is disabled, right-click on the volume again and select **Fast Resync > Add** from the pop-up menu. This enables FastResync.
2. Prepare a new mirror for Instant Recovery backup.

- a. Create a new mirror by right-clicking on the volume and selecting **Snap > Snap Start**.
- b. Make sure FastResync is enabled on the mirror. Click **OK** to create the mirror and start full synchronization.



3. On the **Mirrors** tab, ensure that synchronization has completed as indicated by **Snap Ready** in the **Status** field.



Instant Recovery for Databases

To configure an Instant Recovery policy for database clients, refer to the appropriate NetBackup database agent guide.

Snapshot Configuration Notes

8

This chapter provides notes and configuration assistance for the Advanced Client snapshot methods.

For help using the NetBackup Administration Console to select a snapshot method for a policy, refer to “[Configuring an Advanced Client Policy](#)” on page 79.

The following topics are covered in this chapter:

- ◆ [nbu_snap](#)
- ◆ [VxFS_Checkpoint](#)
- ◆ [VxFS_Snapshot](#)
- ◆ [vxvm](#)
- ◆ [FlashSnap](#)
- ◆ [VVR](#)
- ◆ [NAS_Snapshot](#)
- ◆ [VSS_Transportable](#)
- ◆ [Array-Related Snapshot Methods:](#)

TimeFinder, ShadowImage, BusinessCopy, VSS_Transportable



nbu_snap

The **nbu_snap** snapshot method is for Solaris clients only. It is for making copy-on-write snapshots for UFS or VERITAS VxFS file systems.

The information in this section applies to either Standard or FlashBackup policy types.

Note nbu_snap is not supported in clustered file systems, either as the selected snapshot method, or as the default snapctl driver when configuring FlashBackup in the earlier manner.

An alternative copy-on-write snapshot method, for clustered file systems, is VxFS_Snapshot with a FlashBackup policy.

Note nbu_snap does not support VxVM volumes that belong to a shared disk group.

Cache device

- ◆ The cache device is a raw disk partition: either a logical volume or physical disk. This is used for storing the portions of the client's data that are changed by incoming write requests while the copy-on-write is in progress.
- ◆ For the cache device, do not select an active partition containing valuable data. Any data in that partition will be lost when the snapshot is complete.

Caution Choose a cache partition carefully! The cache partition's contents will be overwritten by the snapshot process.

- ◆ Specify the raw partition as the full path name of either the character special device file or the block device file. For example:

Solaris raw partition:

`/dev/rdsk/c2t0d3s3`

Or

`/dev/dsk/c2t0d3s3`

VxVM volume:

`/dev/vx/rdsk/diskgroup_1/volume_3`

Or

`/dev/vx/dsk/diskgroup_1/volume_3`

Note Do not specify wildcards (such as `/dev/rdisk/c2*`) as paths.

- ◆ The cache partition must be unmounted.
- ◆ The cache partition must reside on the same host as the snapshot source (the client's data to back up).
- ◆ The partition must have enough space to hold all the writes to the partition that may occur during the backup. Note that backups during off-peak periods normally require a smaller cache than a backup during peak activity. (See [“Sizing the Cache Partition”](#) for more suggestions on cache size.)
- ◆ For the Media Server or Third-Party Copy Device method, the host containing the snapshot source and cache must be visible to the media server or third-party copy device (refer to the chapter titled [“SAN Configuration for Advanced Client”](#)).
- ◆ For the Media Server or Third-Party Copy Device method, the disk containing the cache must meet the requirements spelled out under [“Disk Requirements for Media Server/ Third-Party Copy”](#) on page 192.

Sizing the Cache Partition

The size needed for the cache partition depends on user write activity during the backup, not on the size of the client's file system. If the backup occurs when user activity is heavy, a larger cache is required.

Use the following procedure to determine a size for the cache partition:

1. Consider the period in which the backup will occur: the more user activity expected, the larger the cache required.

You should execute the following procedure at an appropriate period, when your snapshot backups typically run. If user activity at your site is known to vary with the time of day, a different time could bring very different results.

2. Make sure a raw partition is available on a separate disk (see [“Cache device”](#) on page 142 for cache partition requirements).
3. During the appropriate backup period, create an nbu_snap snapshot by entering the following as root:

```
/usr/opensv/netbackup/bin/driver/snapon snapshot_source cache
```

where *snapshot_source* is the partition on which the client's file system is mounted, and *cache* is the raw partition to be used as copy-on-write cache. For example:

```
/usr/opensv/netbackup/bin/driver/snapon /omo_cat3  
/dev/vx/rdsk/zeb/cache
```



Example output:

```
matched /omo_cat3 to mnttab entry /omo_cat3
mount device: /dev/vx/dsk/omo/vol03 fstype: vxfs
snapshot 29 enabled on /omo_cat3 at 06/05/03 15:16:02
```

4. In `/usr/opensv/netbackup/bin/driver`, enter the `snaplist` and `snapcachelist` commands.
 - ◆ `snaplist` - shows the id of each snapshot, the size of the partition containing the client file system, and the amount of file system write activity in 512-byte blocks that occurred during the `nbu_snap` snapshot (under the `cached` column). The more blocks cached as a result of user activity, the larger the cache partition required.
 - ◆ `snapcachelist` - shows each cache device in use and what percentage has been used (busy). For each cache device listed, `busy` shows the total space used in the cache. This value indicates the size of the raw partition that may be required for `nbu_snap` cache.

For more details on the snap commands, refer to “[nbu_snap Commands](#)” on page 239.

Note The snap commands can be used in a script.

5. If the cache partition is not large enough, the backup will fail with status code 13, “file read failed.” The `/var/adm/messages` log may contain errors such as the following:

```
Mar 24 01:35:58 bison unix: WARNING: sn_alloccache: cache
/dev/rdsk/c0t2d0s3 full - all snaps using this cache are now
unusable
```
6. Using the information provided by `snaplist` and `snapcachelist`, you have several options:
 - ◆ Specify a larger (or smaller) partition as cache, depending on the results shown by `snaplist` and `snapcachelist`.
 - ◆ Reschedule backups to a period when less user activity is expected.
 - ◆ If multiple backups are using the same cache, reduce the number of concurrent backups by rescheduling some of them.
7. When finished with the snapshot, you can remove it by entering the following:

```
/usr/opensv/netbackup/bin/driver/snapoff snapid
```

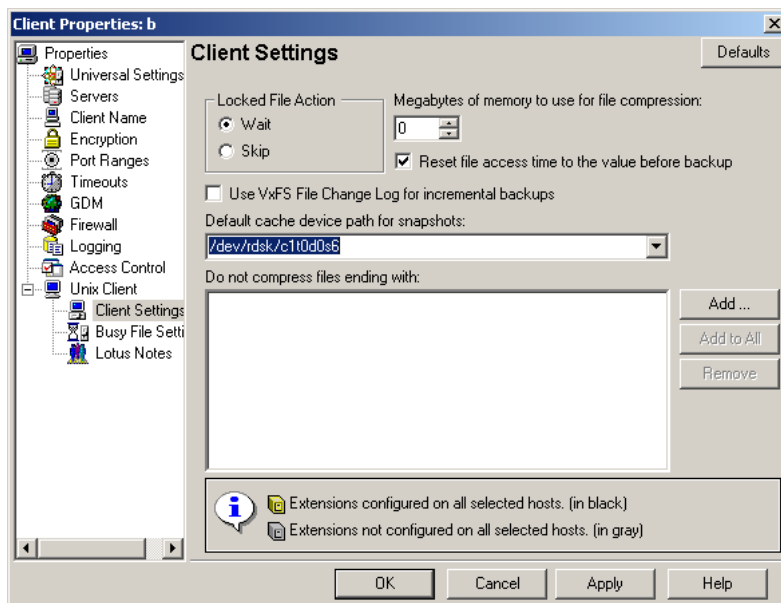
where *snapid* is the numeric id of the snapshot created at [step 3](#).

Note A snapshot created manually with the `snapon` command is not controlled by a NetBackup policy. When run manually, `snapon` creates a copy-on-write snapshot only. The snapshot will remain on the client until it is removed by entering `snapoff` or the client is rebooted.

How to Enter the Cache

For the `nbu_snap` and `VxFS_Snapshot` methods, you must identify a raw partition to be used by the copy-on-write process. Depending on how the policy is configured, this can be done in any of several ways:

- ◆ You can specify the raw partition when using the **Snapshot Policy Configuration wizard** to create an Advanced Client policy.
- ◆ If manually selecting the snapshot method on the Advanced Snapshot Options dialog, you have two options for specifying the raw cache partition:
 - a. Use the **Host Properties > Clients > Client Properties dialog > UNIX Client > Client Settings** to specify the raw partition in the **Default cache device path for snapshots** field. This setting applies to the client in all policies.



- b. Use the Advanced Snapshot Options dialog, **Cache device path Value** field. This cache setting applies to all clients in the current policy, and overrides the cache setting in the Client Settings dialog.

Snapshot Options – Policy test

The following selections are optional. If no snapshot method is selected, NetBackup will select a snapshot method at the time of backup.

Snapshot method for this policy:

nbu_snap

The following parameters can be set for backups performed in this policy that use this snapshot method:

Parameter	Value
Cache device path	

OK Cancel Help

- ◆ If you want NetBackup to be able to select the nbu_snap or VxFS_Snapshot methods by means of the auto method, specify the cache on the **Host Properties > Clients > Client Properties dialog > UNIX Client > Client Settings** as described above.
- ◆ In a FlashBackup policy: if **Perform snapshot backups** is NOT selected, you must use a **CACHE=** directive in the Backup Selections tab. This cache setting applies to all clients in the current policy, and overrides the cache setting in the Host Properties dialog. (This means of configuring cache will be discontinued in a future release.)

VxFS_Checkpoint

The VxFS_Checkpoint snapshot method is for making copy-on-write snapshots. This is one of several snapshot methods that support Instant Recovery backups. Note that for VxFS_Checkpoint, the Instant Recovery snapshot is made on the same disk file system that contains the client's original data.

For VxFS_Checkpoint, VxFS 3.4 or later with the Storage Checkpoints feature must be installed on the NetBackup clients (HP requires VxFS 3.5; AIX and Linux require VxFS 4.0).

- ◆ The VxFS_Checkpoint method is not supported for backing up raw partitions (whether **FlashBackup** or **Standard** policies).
- ◆ Make sure there is enough disk space available for the checkpoint. The file system containing the snapshot source should have at least 10% free space in order to successfully implement the checkpoint.

VxFS Multi-Device System

VxFS_Checkpoint and vxvm are the only snapshot methods in Advanced Client that support the multi-device system (MDS) feature of VxFS 4.0.

The multi-device file system feature requires a VxVM 4.0 volume set. With volume sets, you can group related volumes into one volume set, and mount a VxFS file system on it. This means that a VxFS file system can be mounted on more than one volume. This feature allows file systems to make best use of the different performance and availability characteristics of the underlying volumes. For example, file system metadata could be stored on volumes with higher redundancy, and user data on volumes with better performance.

For background and configuration assistance with multi-device file system, refer to the *VERITAS File System Administrator's Guide* and the *VERITAS Volume Manager 4.0 Administrator's Guide*.

Note Offhost backup is not supported for a VxFS 4.0 multi-device system.

Storage Checkpoint Disk Usage

The `ls` command does not list Storage Checkpoint disk usage. This means that the primary volume may appear to have available space even if it is full. You must use the `fsckptadm list` command to show Storage Checkpoint disk usage. Refer to the *VERITAS File System Administrator's Guide* for more information on `fsckptadm`.



Note A new Storage Checkpoint is created whenever a VxFS_Checkpoint policy is executed.

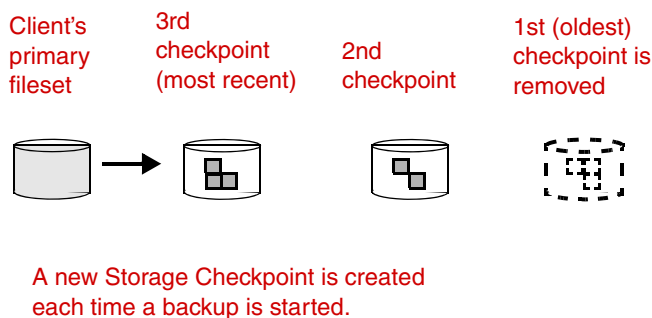
Checkpoint Retention Schedules

For Instant Recovery storage checkpoints, no data is moved between the client and the media server. Instead, a VxFS Storage Checkpoint is created on the client, and, for Oracle clients, the file and directory names are sent to the server for the catalog. In the case of file systems only (non-database clients), only the directory name is sent, not the file data.

Changes made to the files on the client's primary fileset are reflected in the Storage Checkpoint until the backup policy is again executed, creating another Storage Checkpoint. Storage Checkpoints are created and retained until the maximum checkpoints threshold is exceeded, when the oldest checkpoint is removed.

The next figure shows the Instant Recovery checkpoint expiration schedule. If the maximum checkpoint value is set to 2 and a third checkpoint is created, the oldest is removed.

Instant Recovery Retention Schedule for Storage Checkpoints.



Block-Level Restore

If only a small portion of a file system or database changes on a daily basis, full restores are unnecessary. The VxFS Storage Checkpoint mechanism keeps track of data blocks modified since the last checkpoint was taken. Block-level restores take advantage of this by restoring only changed blocks, not the entire file or database. This leads to faster restores when recovering large files.

Refer to [“Instant Recovery: Block-Level Restore \(UNIX Clients Only\)”](#) on page 197 for instructions on setting up this feature.

VxFS_Snapshot

The VxFS_Snapshot method is for making copy-on-write snapshots of local Solaris or HP clients. Offhost backup is not supported with this snapshot method.

Note VxFS_Snapshot supports the FlashBackup policy type only.

Please note the following:

- ◆ The VxFS_Snapshot method can only be used to back up a single file system. If multiple file systems are specified in the policy's Backup Selections list when using this method, the backup will fail.

Note In a FlashBackup policy, if the Backup Selections list is configured with CACHE= entries (see [“Configuring FlashBackup in the Earlier Manner \(UNIX only\)”](#) on page 104), FlashBackup does support the backup of multiple file systems from a single policy. For each file system, a separate cache must be designated with the CACHE= entry.

- ◆ You must designate a raw partition to be used for copy-on-write cache.

Raw partition example:

Solaris:

```
/dev/rdsk/c1t0d0s3
```

Or

```
/dev/dsk/c1t0d0s3
```

HP:

```
/dev/rdsk/c1t0d0
```

Or

```
/dev/dsk/c1t0d0
```

See [“Cache device”](#) on page 142 for general requirements for cache partitions. See also [“How to Enter the Cache”](#) on page 145.

- ◆ VxFS_Snapshot is the default snapshot method for FlashBackup clients running HP, when **Perform snapshot backup** is not selected for the backup policy.



vxvm

The **vxvm** snapshot method is for making mirror snapshots with VERITAS Volume Manager 3.1 or later snapshot mirrors. (On Windows, make sure that VxVM has the latest VxVM service packs and updates.)

The **vxvm** snapshot method works for any file system mounted on a VxVM volume. However, before the backup is performed, the data must be configured with a VxVM 3.1 or later snapshot mirror or a VxVM 4.0 or later cache object (otherwise, the backup will fail).

- ◆ For help configuring a snapshot mirror, refer to “[Creating a Snapshot Mirror of the Source](#),” below, or refer to your *VERITAS Volume Manager* documentation.
- ◆ For help configuring a cache object, refer to your *VERITAS Volume Manager* documentation, and to “[VxVM Instant Snapshots](#)” on page 151.
- ◆ For Instant Recovery backups of data configured on VxVM volumes on Windows, the VxVM volume names must be 12 characters or less. Otherwise, the backup will fail.

Note vxvm and VxFS_Checkpoint are the only snapshot methods in Advanced Client that support the multi-device system (MDS) feature of VxFS 4.0.

Note Since VxVM does not support fast mirror resynchronization on RAID 5 volumes, the **vxvm** snapshot method must not be used with VxVM volumes configured as RAID 5. If the **vxvm** snapshot method is selected for a RAID 5 volume, the backup will fail.

Creating a Snapshot Mirror of the Source

To use the **vxvm** snapshot method with VxVM volumes in the third-mirror (split-mirror) configuration, you must create a snapshot mirror of the snapshot source before making a backup. Before specifying a snapshot source as described in this chapter, create a snapshot mirror on the client by using one of the following methods:

- ◆ In the Volume Manager Storage Administrator interface:

UNIX: select the volume that will be the snapshot source, right click on it, and select **Snapshot** from the pop-up menu. In the Volume Snapshot dialog, select **Enable FMR** (if available, see note below) and click the **Snapstart** button.

Windows: select the volume that will be the snapshot source, right click on it, select **Snap**, and choose **Snapstart**.

For details, refer to your VERITAS Volume Manager documentation.

Or (for UNIX):

- ◆ Enter the following commands:

```
/usr/sbin/vxassist -g disk_group snapstart volume_name  
/usr/sbin/vxvol -g disk_group set fmr=on volume_name
```

where:

- ◆ *disk_group* is the Volume Manager disk group to which the volume (snapshot source) belongs.
- ◆ *volume_name* is the name of the volume designated at the end of the snapshot source path (for example, *vol1* in */dev/vx/rdisk/dg/vol1*).
- ◆ *fmr=on* sets the Fast Mirror Resynchronization attribute, which resynchronizes the mirror with its primary volume but only copies the blocks that have changed, rather than performing a full resynchronization. Fast mirror resynchronization can dramatically reduce the time required to complete the backup.

Fast Mirror Resynchronization (FMR) is a separate product for VERITAS Volume Manager.

Note If the Media Server or Third-Party Copy method is used, the disks that make up the disk group must meet the requirements spelled out under “[Disk Requirements for Media Server/ Third-Party Copy](#)” on page 192.

VxVM Instant Snapshots

Advanced Client supports two additional kinds of snapshot volumes included in Volume Manager 4.0: full-sized instant snapshots and space-optimized instant snapshots. These offer advantages over traditional third-mirror snapshots, such as immediate availability and easier configuration and administration.

- ◆ **Full-sized instant snapshot**

This is a variation on the VxVM third-mirror volume snapshot model. It makes a snapshot volume available for access as soon as the snapshot plexes have been created. Like traditional third-mirror volumes, this volume, after resynchronization, can be moved into a separate disk group or turned into an independent volume.

- ◆ **Space-optimized instant snapshot**

This type of snapshot volume contains only the blocks that changed during the snapshot, and uses a storage cache (cache object) to store them. The size of this cache may be configured when the snapshot is created. This volume can be created very quickly and uses a minimum of disk space. Note that it cannot be moved into a separate disk group or turned into an independent volume.



Refer to *VERITAS Volume Manager 4.0 Administrator's Guide* for complete descriptions of instant snapshot volumes and for configuration assistance.

NetBackup Snapshot Methods

The instant snapshot feature of VxVM 4.0 is supported by NetBackup's vxvm, FlashSnap, and VVR methods. Please note:

- ◆ FlashSnap supports full-sized instant snapshots only.
- ◆ For alternate client backup, only the VVR method supports space-optimized snapshots.

Apart from configuring the VxVM volumes and selecting vxvm, FlashSnap, or VVR as the NetBackup snapshot method, there are no special parameters in NetBackup to be specified.

Space-Optimized Instant Snapshots

When using the space-optimized snapshot feature of VxVM, you must create a cache object for the snapshot. For the procedure, refer to "[Creating space-optimized snapshots](#)" on page 136.

FlashSnap

FlashSnap uses the Persistent FastResync and Disk Group Split and Join features of VERITAS Volume Manager (VxVM).

The FlashSnap snapshot method can be used for alternate client backups only, in the split mirror configuration, which is described under “[Split mirror](#)” on page 11.

Note FlashSnap supports VxVM full-sized instant snapshots but not space-optimized snapshots. For more information, refer to “[VxVM Instant Snapshots](#)” on page 151.

Note FlashSnap does not support VxVM volumes that belong to a shared disk group.

Testing Volumes for FlashSnap

Before running an alternate client backup with the FlashSnap snapshot method, you should test your FlashSnap volume configuration as described below. The goal of this procedure is to ensure that the disk(s) containing the volume can be moved (deported and imported) to the alternate client without errors. For instance, if the disks containing the backup volumes contain some part of a volume that has another part on a disk not in the backup disks, then the backup disks cannot be split off into a new disk group and deported.¹

On UNIX

The following steps are described in more detail in the *VERITAS FlashSnap Point-In-Time Copy Solutions Administrator's Guide*.

1. On the primary host:

- a. If not already done, create a snapshot mirror:

```
vxassist -g diskgroup -b snapstart volume
```

- b. Add a DCO log to the volume:

```
vxassist -g diskgroup addlog volume logtype=dco
```

- c. Enable FastResync on the volume:

```
vxvol -g diskgroup set fastresync=on volume
```

- d. Create a snapshot volume from the primary volume:

1. Deporting a disk group means disabling access to that disk group. See the *Volume Manager Administrator's Guide* for more information on deporting disk groups.



```
vxassist -g diskgroup snapshot volume snap_volume
```

- e. Move the disks containing the snapshot volume to a separate (split) disk group:

```
vxvg split diskgroup split_diskgroup snap_volume
```

If the volume has not been properly configured, you may see an error similar to the following:

```
host-name# vxvg split lhdvvr lhdvvr_split SNAP-emc_concat
vxvm:vxvg: ERROR: vxvg split lhdvvr lhdvvr_split failed
vxvm:vxvg: ERROR: emc_dis05 : Disk moving, but not all
subdisks on it
```

Look again at the layout of the disks and the volumes assigned to them, and reassign the unwanted volumes to other disks as needed. Consult the *VERITAS FlashSnap Point-In-Time Copy Solutions Administrator's Guide* for examples of disk groups that can and cannot be split.

- f. Deport the split disk group:

```
vxvg deport split_diskgroup
```

2. On the secondary host:

- a. Import the disk group that was deported from the primary:

```
vxvg import split_diskgroup
```

- b. Enable the imported volume:

```
vxvg recover -g split_diskgroup -m snap_volume
```

- c. Start the volume:

```
vxvol -g split_diskgroup start snap_volume
```

If the above commands execute successfully, the volume setup is correct.

Note After doing this test, you must re-establish the original configuration to what it was prior to testing the volumes. 1, deport the disk group on the alternate client, then 2, import the disk group on the primary client, and 3, recover and join the original volume group. For directions, refer to “[For FlashSnap \(Solaris, HP, AIX, Linux\):](#)” on page 224.

On Windows

1. On the primary host:

- a.** If not already done, create a snapshot mirror:

```
vxassist snapstart \Device\HarddiskDmVolumes\diskgroup\volume
```

- b.** Create a snapshot volume from the primary volume:

```
vxassist snapshot \Device\HarddiskDmVolumes\diskgroup\volume  
DrivePath=C:\Temp\Mount SNAP-Volume
```

- c.** Move the disks containing the snapshot volume to a separate (split) disk group.

Disk group is also deported after this command completes:

```
vxdbg -g DskGrp -n SPLIT-DskGrp split  
\Device\HarddiskDmVolumes\diskgroup\snap_volume
```

2. On the secondary host:

- a.** Rescan to make the deported disk group visible on the secondary host:

```
vxassist rescan
```

- b.** Import the disk group that was deported from the primary:

```
vxdbg -g split_diskgroup import
```

- c.** Assign the snapshot volume to an empty NTFS directory.

This example uses C:\Temp\Mount.

```
vxassist assign  
\Device\HarddiskDmVolumes\split_diskgroup \snap_volume  
DrivePath=C:\Temp\Mount
```

If the above commands execute successfully, the volume setup is correct.



VVR

The VVR snapshot method (for UNIX clients only) relies on the VERITAS Volume Replicator, which is a licensed component of VxVM. The Volume Replicator maintains a consistent copy of data at a remote site. Volume Replicator is described in the *VERITAS Volume Replicator Administrator's Guide*.

The VVR snapshot method can be used for alternate client backups only, in the data replication configuration, described under “[Data Replication \(UNIX Only\)](#)” on page 14. VVR makes use of the VxVM remote replication feature. The backup processing is done by the alternate client at the replication site, not by the primary host or client.

Note VVR supports VxVM instant snapshots. For more information, refer to “[VxVM Instant Snapshots](#)” on page 151.

Set Up Volume Replication

Before doing a replication backup with VVR, make sure to configure the Volume Replicator as explained in the *Volume Replicator Administrator's Guide*.

Name Registration

Inband Control (IBC) messages are used to exchange control information between primary and secondary hosts. A name has to be registered at both primary and secondary host for each replicated volume group before IBC messaging can be used. The VVR snapshot method assumes that the application name is APP_NBU_VVR. To avoid an initial backup failure, you should register that name as described in [step 1](#) on page 157.

Note If APP_NBU_VVR is not registered, NetBackup will register the name when the first backup is attempted, but the backup will fail. Subsequent backups, however, will succeed.

Primary/Secondary Disk Group and Volume Names

For the VVR snapshot method, the disk group and volume must have the same name on both the primary and secondary host. If the names are different, the VVR backup will fail.

Test the Replication Setup

Before running an alternate client backup with VVR, you should test your replication setup as follows.

1. On both primary and secondary host, register the APP_NBU_VVR name:

```
vxibc -g diskgroup register APP_NBU_VVR replicated_group
```

The above command must be executed twice, once on each host.

2. On the primary host, send an IBC message to the secondary host:

```
vxibc -g diskgroup send APP_NBU_VVR replicated_group  
replication_link
```

3. On the secondary host, receive the IBC message from the primary host:

```
vxibc -g diskgroup -R10 receive APP_NBU_VVR replicated_group
```

4. On the secondary host, restart replication:

```
vxibc -g diskgroup unfreeze APP_NBU_VVR replicated_group
```

If the above commands execute successfully, the replication setup is correct.



NAS_Snapshot

NetBackup can make point-in-time snapshots of data on NAS (NDMP) hosts using the NDMP V4 snapshot extension. The snapshot is stored on the same device that contains the NAS client data. From the snapshot, you can restore individual files or roll back a file system or volume by means of the Instant Recovery feature.

Note NetBackup for NDMP software is required on the server, and the NAS vendor must support the NDMP V4 snapshot extension.

Note The backup retention period specified in Master Server Properties applies only to Network Appliance SnapVault snapshots using a SnapVault storage unit. If the snapshot is not configured with a SnapVault storage unit, the only means of controlling snapshot deletion is the **Maximum Snapshots (Instant Recovery Only)** parameter on the Advanced Snapshot Options dialog of the policy. Refer to [“Configuring Snapshot Deletion”](#) on page 128 for more information.

For help in setting up a policy for NAS snapshot, see the [“NAS Snapshot Configuration”](#) chapter.

VSS_Transportable

This method uses the Volume Shadow Copy Service of Windows 2003. It is for alternate client backups of Windows 2003 clients, where the client data is stored on either a disk array such as EMC or Hitachi, or in a VERITAS Storage Foundation for Windows 4.1 or later volume. VSS_Transportable supports file system backup of a disk partition (such as E:\) and backup of Exchange databases.

Note VSS-based snapshot methods offer a general interface to Windows Shadow Copy Services. The actual snapshot method used for the backup depends on the VSS snapshot providers available on your system. When a VSS-based snapshot method is configured in the NetBackup policy, Windows Shadow Copy Services selects the actual snapshot mechanism.

For a list of arrays supported as of this printing, see “[Arrays Currently Supported for VSS_Transportable](#)” on page 163. For the most up-to-date list of disk arrays supported by this method, see the *NetBackup Advanced Client Configuration and Compatibility* document available on the VERITAS support site (see the preface of this manual for help accessing that document).

Prerequisites

- ◆ *If client data is stored on a disk array such as EMC or Hitachi:* you must set up primary and secondary (mirror) disks, and client data must be visible to the alternate client on a mirrored or secondary volume only, configured as a split mirror. See “[Configuring Primary and Secondary Disks](#)” on page 167. Assistance from the array vendor may be required.
- ◆ *If client data is stored in a Storage Foundation for Windows volume (VxVM):* before running an alternate client backup with the VSS_Transportable snapshot method, configure and test volumes as described under “[Testing Volumes for FlashSnap](#)” on page 153.

Notes and Restrictions on VSS_Transportable

- ◆ Supports backup of Storage Foundation for Windows 4.1 or later logical volumes *OR* backup of file systems on supported disk arrays. Backup of VxVM volumes configured on disk arrays is not supported.
- ◆ Supports backup of Windows NTFS file systems on a disk partition, and backup of data in an Exchange database. The policy type can be MS-Exchange-Server or MS-Windows-NT.



- ◆ The mirror volumes must be configured so as to be visible only to the secondary (alternate) client. If both primary and mirror volumes are visible to the same client, two disks with identical file system and volume naming are presented to the client when the mirror is split. A conflict results and the backup fails.
- ◆ Does not support Instant Recovery.
- ◆ Does not support the backup of Windows system-protected files (the System State, such as the Registry and Active Directory). If the volume containing the data to back up includes Windows system files, that volume cannot be backed up with the VSS_Portable snapshot method.
- ◆ Does not support the backup of Windows system database files (such as RSM Database and Terminal Services Database).

Array-Related Snapshot Methods

This section describes the snapshot methods that are designed for particular hardware arrays (TimeFinder, ShadowImage, BusinessCopy). Also included are notes on VSS_Transportable (see [“Configuring Primary and Secondary Disks”](#) on page 167).

Configuration Checklist

This checklist includes major caveats and important information. READ THIS TABLE before setting up your disk arrays for the array-specific snapshot methods. The right column refers to sources of further information.

CHECK THE FOLLOWING !	Refer to these topics for help
If you want your client data configured over Volume Manager volumes, make sure your arrays and operating system are supported by Volume Manager (VxVM).	Refer to the <i>NetBackup Release Notes</i> , or see “Advanced Client Assistance” on page xvi.
Make sure the client data is correctly mirrored to secondary disks in the array.	See “Configuring Primary and Secondary Disks” on page 167.
When configuring a backup policy, be sure to select a snapshot method that supports your arrays.	See “The Snapshot Methods” on page 162.
For NetBackup Media Server or Third-Party Copy Device offhost methods: ask your array support technician to configure your array as follows: <ul style="list-style-type: none"> • The NetBackup clients must have access to primary and secondary disks in the array. • The media server must have access to the secondary disks. 	See “Disk Configuration Requirements” on page 165.
Solaris: If client data is configured over Volume Manager volumes, label all secondary disks using the <code>format</code> command (<code>label</code> option).	See “Disk Configuration Requirements” on page 165.
Solaris: The EMC Symmetrix/DMX array must be configured in <i>Common Serial Number Mode</i> to support multiple client SCSI and/or fibre channel connections.	See “Multiple Connectivity to EMC Array: Common Serial Number mode” on page 167
Do not include secondary disks in a Volume Manager disk group. Be sure to follow this and other restrictions when using Volume Manager.	See “Disk Types” on page 182.



CHECK THE FOLLOWING !	Refer to these topics for help
Read the “ Best Practices ” section.	See “ Best Practices ” on page 186.

Overview

This section describes the array-related snapshot methods provided in Advanced Client, explains the need for data mirroring, and introduces terms used in this chapter.

The Snapshot Methods

For mirror-type snapshot backups on supported disk arrays, select the TimeFinder, ShadowImage, or BusinessCopy snapshot method according to the vendor-type of disk array that contains the client data. If the snapshot method does not match the vendor-type of the array, the backup will fail.

VSS_Transportable, on the other hand, is not designed for any particular array.

◆ TimeFinder

The **TimeFinder** snapshot method is for making mirror snapshots on EMC Symmetrix/DMX disk arrays with TimeFinder SYMCLI (with or without VERITAS Volume Manager 3.1 or later). Supports UFS and VxFS file systems, VxVM logical volumes, and raw partitions.

◆ ShadowImage

The **ShadowImage** snapshot method is for making mirror snapshots on Hitachi Data Systems (HDS) disk arrays with ShadowImage (HOMRCF). Supports UFS and VxFS file systems, VxVM logical volumes, and raw partitions.

◆ BusinessCopy

The **BusinessCopy** snapshot method is for making mirror snapshots on HP XP series disk arrays with BusinessCopy Services. Supports UFS and VxFS file systems, VxVM logical volumes, and raw partitions.

◆ VSS_Transportable

For a description, see “[VSS_Transportable](#)” on page 159. This method is not designed for a particular array type.

Each of the following methods must be used for its own array-type

Match Array-Based Snapshot Method to Type of Array

To back up the following:	Use this snapshot method:
EMC Symmetrix/DMX disk arrays	TimeFinder
Hitachi disk arrays	ShadowImage
HP XP disk arrays	BusinessCopy
<ul style="list-style-type: none"> ◆ These snapshot methods cannot be switched: selecting TimeFinder to back up an Hitachi array will cause the backup to fail. ◆ For the latest information on supported disk arrays, see “Advanced Client Assistance” on page xvi. 	

Configuration of client data over VxVM volumes is supported only on certain combinations of disk arrays and platforms. You can go to the VERITAS support web site to locate an up-do-date list of VxVM supported arrays/platforms (see “[Advanced Client Assistance](#)” on page xvi).

If client data is not configured over VxVM, all above arrays are supported (no restrictions).

Note As an alternative, the vxvm snapshot method can be used in backing up any of the above disk arrays, if the client data is configured over Volume Manager volumes.


Arrays Currently Supported for VSS_Transportable

Arrays supported for VSS_Transportable
EMC CLARiiON
Hitachi
<ul style="list-style-type: none"> ◆ For the latest information on supported disk arrays, see “Advanced Client Assistance” on page xvi.




Client Data Must Be Mirrored

When NetBackup makes a mirror-type snapshot, the client data on the primary disk must be mirrored on a secondary disk prior to the backup.



If the client's data is distributed across two or more primary disks by means of a VxVM volume, an equal number of mirror disks must also contain the same data.

NetBackup client data in VxVM volumes:
where vol_1 is distributed across primary disk_1 and disk_2



Disk Terms

The terms used in this manual for array disk mirroring are *primary* and *mirror* (or *primary* and *secondary*). Some array vendors refer to these as follows:

- ◆ EMC: The primary is called the *standard*, and the mirror is called a *BCV*.

- ◆ Hitachi and HP: Primary and secondary are called *primary volume* and *secondary volume*.

Disk Configuration Requirements

Contact the array's support technicians for help in configuring arrays to your specifications.

Access to Disk Arrays


- ◆ For NetBackup Media Server or Third-Party Copy Device methods only: NetBackup clients must have access to both primary and secondary disks (via SCSI or fibre channel, or both). If the clients do not have access to both primary and secondary disks, the backup will fail.
- ◆ For NetBackup Media Server or Third-Party Copy Device methods only: the media server requires access to the secondary disks only (via SCSI or fibre channel, or both). If the media server does not have access to the secondary disks, the backup will fail.

Note Although the above configuration is required by NetBackup, a support technician for your disk array vendor must configure this access.


Connection to Disk Array: SCSI and Fibre Channel

NetBackup supports three configurations, each requiring setup assistance from your array vendor. Note that Fibre Channel and SCSI are both supported.


Configuration 1: Local SCSI (no Fibre Channel)



Configuration 2: Array on Fibre Channel



Configuration 3: Array Behind Bridge; Bridge on Fibre Channel



Multiple Connectivity to EMC Array: Common Serial Number mode

EMC Symmetrix and EMC DMX disk arrays with multiple connections from NetBackup clients and media servers (multiple SCSI or both fibre channel and SCSI) must be configured in Common Serial Number Mode.

If the arrays are not configured in Common Serial Number Mode, and there are multiple SCSI connections (or fibre channel and SCSI connections) to the same disks in the array, NetBackup and Volume Manager will be presented with two different serial numbers for the same disk: one for the SCSI path and one for the fibre channel path. As a result:

- ◆ Volume Manager will be confused and will not be able to run in DMP (Dynamic Multipath) mode.
- ◆ NetBackup data movement services will be confused, since the disk serial number is used to identify the proper disk to back up.

To prevent these problems, the array must be configured in Common Serial Number Mode.

Caution If Common Serial Number Mode is not configured for an EMC disk array that has multiple client and media server connections, the backup may fail.

Configuring Primary and Secondary Disks

The following is a brief listing of the commands required for making the primary-to-mirror disk association, for EMC, Hitachi, and HP arrays. The primary-to-mirror association can be set up before or after installation of NetBackup Advanced Client, but must be done prior to running the backup. The primary-to-mirror association is done on the NetBackup client only.

Note If a mirror disk is not correctly associated and synchronized with the primary disk, a snapshot of the client's data cannot be made. (A snapshot has to be made on the mirror, not on the primary.) In that case, if the backup policy is configured with a mirror-type snapshot, the backup will fail.

EMC CLARiiON

For an EMC CLARiiON disk array on the NetBackup client, you must create device groups, add primary and mirror (secondary) devices to the groups, and associate or pair the primaries with the secondaries. Once associated, the secondary disks must be



synchronized with the primary disks. During synchronization, the primary disks are copied to the secondaries. Please refer to your EMC Navisphere Host Based Agent/CLI documentation for instructions.

EMC Symmetrix / DMX

For an EMC Symmetrix or DMX disk array on the NetBackup client, you must create device groups, add primary and mirror (secondary) devices to the groups, and associate or pair the primaries with the secondaries. Once associated, the secondary disks must be synchronized with the primary disks. During synchronization, the primary disks are copied to the secondaries.

Use the following commands.

Note Please refer to your EMC TimeFinder SYMCLI documentation for more details on these commands.

symdg

Creates a disk group.

symld

Adds primary disk to the disk group.

symbcv

Associates a secondary disk with a primary disk.

▼ Create the EMC disk groups

1. Create a disk group that will contain any number of primary and secondary disks.

```
symdg create nbfim_test
```

Creates disk group named nbfim_test.

2. Add primary disks to the disk group.

```
symld -g nbfim_test add dev 02A
```

Adds primary disk 02A to disk group nbfim_test.

3. Add secondary (BCV) disks to the disk group.

```
symbcv -g nbfim_test associate dev 08C BCV001
```


Adds the secondary disk 08C to the same disk group.

4. Synchronize the secondary disks with the primaries in the disk group.

```
symmir -g nbfim_test establish
```

Pairs, or associates, the primary with the mirror, and synchronizes the mirror with the primary. If there are multiple primaries and mirrors, they are paired according to the order in which they were added to the group.

5. Show the output.

```
symmir -g nbfim_test query
```

When the above commands are successfully entered, NetBackup can execute snapshot requests involving primary device 02A and its associated mirror 08C.

Hitachi and HP Arrays

The procedure for setting up Hitachi arrays is identical to that for HP arrays. For more detail on the commands and files described below, refer to your Hitachi Data Systems or HP documentation.

The basic steps are these (details follow), entered on the NetBackup client only:

1. Create array configuration files.
2. Add array service names to `/etc/services` file.
3. Start the RAID Manager daemons.
4. Set the instance number and enable the listing of output.
5. View the state of the arrays.
6. Configure the arrays, depending on your requirements.

▼ Create array configuration files

You need a configuration file for each set of primary disks, and another file for each set of mirror (secondary) disks. The entries in the file must be space-delimited.

1. Create a configuration file for your primary disks. Use this path and file name:

```
UNIX  
  
/etc/horcmX.conf
```



Windows (for VSS_Transportable)

`%WINDIR%\horcmX.conf`

where *X* is an integer. For example: `/etc/horcm0.conf`. This integer is called the *instance number*.

2. Create a configuration file for your mirror disks, using the same path and file name as above, but with a different instance number.

For example: `/etc/horcm1.conf`

On the following pages are two example files. Note that entries must be separated by spaces.

Except for comment lines (#), the file must contain the `HORCM_MON`, `HORCM_CMD`, `HORCM_DEV`, and `HORCM_INST` parameters, followed by appropriate entries (explained below).

Example 1: Primary Disks

Example 1: configuration file `horcm0.conf` for three *primary* disks

`/etc/horcm0.conf` (UNIX), `%WINDIR%\horcm0.conf` (Windows)

Entries must be space delimited

	HORCM_MON						
	#host	service	poll(10ms)	timeout(10ms)			
Host where the configuration file resides. →	turnip	horcmgr0	1000	3000			
Port name for this instance. →							
	HORCM_CMD						
	#cmd_dev_file	cmd_dev_file	cmd_dev_file				
	/dev/rdisk/c2t8d14s2						
	HORCM_DEV						
	#dev_group	dev_name	port#	TargetID	LU#	MU#	
One line per primary disk. →	wiltest	dev1	CL1-A	8	0		
	wiltest	dev2	CL1-A	8	1		
	wiltest	dev3	CL1-A	8	2		
	HORCM_INST						
	#dev_group	partner host	partner service				
Also enter this in /etc/services file →	wiltest	turnip	horcmgr1				

HORCM_MON

Enter values for the following:

- ◆ **host**: the NetBackup client where this configuration file resides. The NetBackup client accesses the disks specified in this file under **HORCM_DEV**, when backing up or restoring data using the **ShadowImage**, **VSS_Transportable**, or **BusinessCopy** snapshot method.
- ◆ **service**: the port name of the RAID Manager instance (for this configuration file) to be registered in the `/etc/services` file (UNIX only).
- ◆ **poll**: the interval at which the disks are monitored, expressed as tens of milliseconds.



- ◆ `timeout`: time-out period for attempting to communicate with the “partner” service, expressed as tens of milliseconds.

HORCM_CMD

Enter values for the following:

- ◆ `cmd_dev_file`: the command device file(s) for the array. For example:

UNIX

`/dev/rdisk/c2t8d14s2`

Windows (VSS_Transportable)

`\\.\PhysicalDrive4`

Following applies to UNIX only:

You can use the NetBackup `bptpcinfo` command to determine the command device file, as follows:

```
bptpcinfo -d /dev/rdisk -o- | grep CM
```

Below is sample output showing a command device file for an Hitachi device and for an HP device.

Command device files	<code>p=/dev/rdisk/c2t8d14s2</code>	<code>s=HITACHI:OPEN-9-CM:60159001C00</code>
(note “CM”):	<code>p=/dev/rdisk/c2t5d35s2</code>	<code>s=HP:OPEN-3-CM:30436002500</code>

The format of the output is:

```
p=/dev/rdisk/c#t#d#s2 s=VID:PID:SN
```

where:

- ◆ VID (vendor ID) must be HP or HITACHI.
- ◆ PID (product ID) must include -CM.
- ◆ The first five characters of the serial number (SN) must match the serial number of the disks.

In this UNIX example, the command device file for the Hitachi array is `/dev/rdisk/c2t8d14s2` and for the HP array it is `/dev/rdisk/c2t5d35s2`.

Following applies to Windows only, for VSS_Transportable:

You can use the `in RAID` command to determine the command device file, as follows:

```
in RAID -CLI $Phys
```

Below is sample output showing a command device file.

DEVICE_FILE	PORT	SERIAL	LDEV	CTG	H/M/12	SSID	R:Group	PRODUCT_ID
Harddisk13	-	-	-	-	-	-	-	ATLAS V 18 SCA
Harddisk14	-	-	-	-	-	-	-	ST318404LC
Harddisk0	CL2-C	20461	82	-	-	-	-	OPEN-3-CM
Harddisk1	CL2-C	20461	564	-	s/P/ss	9972	5:01-04	OPEN-9
Harddisk2	CL2-C	20461	565	-	s/P/ss	9972	5:01-04	OPEN-9
Harddisk3	CL2-C	20461	566	-	s/P/ss	9972	5:01-04	OPEN-9

HORCM_DEV

Enter values for the following:

- ◆ `dev_group`: a user-defined name of a logical grouping of primary and secondary disks.
- ◆ `dev_name`: a user-defined name assigned to a primary-secondary pair of disks within the logical group.

The `dev_group` and `dev_name` parameters are used on the “pair” configuration commands described later in this section.

- ◆ `port #`: the port number specified for the disk, configured by means of the array’s dedicated console (not from a NetBackup host).
- ◆ `Target ID`: the SCSI or fibre channel target ID number of the disk, configured by means of the array’s dedicated console (not from a NetBackup host).
- ◆ `LUN`: the SCSI or fibre channel logical unit number of the disk, configured by means of the array’s dedicated console (not from a NetBackup host).
- ◆ `MU`: a numeric mirror descriptor for cascading disks (default 0). If you are not using cascading disks, this value may be left blank. A cascading disk has more than one mirror (secondary) associated with a given primary.

HORCM_INST

Enter values for the following:

- ◆ `dev_group`: same as under `HORCM_DEV`.
- ◆ `partner host`: the host where the corresponding secondary (or primary) configuration file resides (may be the same as the host specified under `HORCM_MON`). For this example, the host and partner host are both `turnip`. (See under “partner service” for a discussion of *partner*.)



- ◆ `partner service`: the port name of the RAID Manager instance for the corresponding secondary (or primary) configuration file, to be registered in the `/etc/services` file (UNIX only).

For the example `horcm0.conf` file (`/etc/horcm0.conf` on UNIX, `%WINDIR%\horcm0.conf` on Windows), the `partner service` for `horcmgr0` (entered under `HORCM_MON, service`) is `horcmgr1`. For the secondary-disk configuration example `horcm1.conf` file (below), the `partner service` is the opposite: `horcmgr0`.

Partner is a relative term. From the viewpoint of the configuration file for the primary disks (`horcm0.conf` file), the *partner* file would be `horcm1.conf` (for the secondary disks). It is the same with `partner service` and `partner host`: each refers to the secondary from the viewpoint of the primary, or to the primary from the viewpoint of the secondary.

Note The `partner service` value must be entered in the `/etc/services` file.

Example 2: Secondary Disks

Example 2: configuration file `horcm1.conf`, for three *secondary* disks
`/etc/horcm1.conf` (UNIX), `%WINDIR%\horcm1.conf` (Windows)

Entries must
be space
delimited

Contains same
parameters
(`HORCM_MON`, etc.)
as in config file for
primary disks.
Disk-related entries
refer to the
secondary disks.

```
HORCM_MON
#host      service      poll(10ms)  timeout(10ms)
turnip     horcmgr1      1000        3000

HORCM_CMD
#cmd_dev_file  cmd_dev_file  cmd_dev_file
/dev/rdsd/c2t8d14s2 (UNIX)  OR  \\.\PhysicalDrive0 (Windows)
```

Port name for this
instance

```
HORCM_DEV
#dev_group  dev_name  port#      TargetID  LU#  MU#
wiltest     dev1      CL2-A      16        32
wiltest     dev2      CL2-A      16        33
wiltest     dev3      CL2-A      16        34
```

One line per
secondary
disk

```
HORCM_INST
#dev_group partner host  partner service
wiltest   turnip      horcmgr0
```

Also enter this in
`/etc/services` file
(UNIX only)

See under “[Example 1: Primary Disks](#)” on page 171 (`horcm0.conf`) for a description of these entries.

▼ **Add array service names to `/etc/services` file (UNIX only)**

The values listed under “service” in the configuration files (`horcmgr1` and `horcmgr0` in the above examples) must be entered in `/etc/services` file.

▼ **Restart the `inetd` daemon (UNIX only)**

For example:



```
kill -SIGHUP pid_of_inetd
```

▼ Start the RAID Manager daemons

Enter the following command to start the RAID Manager daemons:

UNIX

```
/bin/horcmstart.sh x x
```

Windows (VSS_Transportable)

```
C:\HORCM\etc\horcmstart x
```

where x is the instance number of each configuration file. For the above UNIX example, the command would be:

```
/bin/horcmstart.sh 0 1
```

The daemons must be running in order to configure your primary and secondary disks.

▼ Set the instance number and enable the listing of output

UNIX

If you are using the Bourne shell, and the instance number for your primary disks is 0, enter the following:

```
HORCMINST=0
HORCC_MRCF=1
export HORCMINST HORCC_MRCF
```

If using the C shell, enter the following:

```
setenv HORCMINST 0
setenv HORCC_MRCF 1
```

Windows

Go to **Control Panel > System > Advanced** and click **Environment Variables** to enter the following variables:

Windows Environment Variables (for VSS_Transportable)

Variable	Value
HORCMINST	0
VSHTCHORCMINST_LOCAL	0
VSHTCHOMRCF_MUN	0

Windows Environment Variables (for VSS_Transportable)

Variable	Value
HORCC_MRCF	1

The HORCMINST parameter determines three things:

- ◆ The array to which commands will be sent.
- ◆ Which disk is the primary and which is the secondary, when using the `paircreate` command (described below).
- ◆ Which disk (primary or secondary) is listed first in each pair when using the `pairdisplay` command to view the state of the arrays (described below). In this example (`HORCMINST=0`), the primaries are listed first. That is because the configuration file that defines the primary disks is named `horcm0.conf`, with 0 as the instance number.

▼ View the state of the arrays

1. To display status information on all the disks, enter the following:

```
pairdisplay -g groupname -CLI -fc
```

where *groupname* is the name specified in the configuration files under `dev_group`. CLI and fc are options:

-CLI formats headers and columns in the resulting display.

-fc includes the percentage of synchronization progress in the display.

For example:

```
pairdisplay -g wiltest -CLI -fc
```

Resulting output:

Group	PairVol	L/R	Port#	TID	LU-M	Seq#	LDEV#	P/S	Status	%	P-LDEV#	M
wiltest	dev1	L	CL1-A	8	0 0	60159	0	P-VOL	PAIR	100	43	-
wiltest	dev1	R	CL2-A	16	32 0	60159	43	S-VOL	PAIR	100	0	-
wiltest	dev2	L	CL1-A	8	1 0	60159	1	P-VOL	PSUS	99	44	W
wiltest	dev2	R	CL2-A	16	33 0	60159	44	S-VOL	SSUS	99	1	-
wiltest	dev3	L	CL1-A	8	2 0	60159	2	SMPL	-	-	-	-
wiltest	dev3	R	CL2-A	16	34 0	60159	45	SMPL	-	-	-	-

2. For status information on a particular pair of disks, enter the following:

```
pairdisplay -g groupname -d dev_name [-CLI] [-fc]
```

where *dev_name* is the name specified in the configuration files under `dev_name`.



Note If no primary-secondary associations (pairings) exist, all disks are listed as SMPL in the P/S column. To create a primary-secondary pairing, see [“If disks are not paired:”](#) on page 180.

The following describes important headers in the **pairedisplay** listing.

Group

This is the `dev_group` name defined in the configuration file.

PairVol

Lists the devices by device name. In the above output, `dev1` is listed twice: the first line is the primary disk, the second is the mirror (secondary). This is shown under the P/S column: P-VOL indicates the primary, S-VOL the secondary.

L/R

Indicates local or remote host, with respect to the current instance number.

Port#

The port number for the disk, configured by means of the array’s dedicated console (not from a NetBackup host).

TID

The SCSI or fibre channel target ID number of the disk, configured by means of the array’s dedicated console (not from a NetBackup host).

LU-M

LU indicates the SCSI or fibre channel logical unit number of the disk, configured by means of the array’s dedicated console (not from a NetBackup host). M is the numeric mirror descriptor for cascading disks. A cascading disk has more than one mirror (secondary) associated with a given primary.

Seq#

This is the unit serial number of the array.

LDEV#

Logical device number of the disk.

P/S

Indicates whether or not the disk is configured in a primary-secondary pair:

- ◆ P-VOL: the disk is the primary.
- ◆ S-VOL: the disk is the secondary.
- ◆ SMPL: the disk is not paired (associated) with any other disk.

Status

Shows the current state of each disk in the array:

- ◆ PAIR: the secondary disk in the pair is synchronized with the primary.
- ◆ PSUS: the pair is split (primary disk).
- ◆ SSUS: the pair is split (secondary disk).
- ◆ COPY: a synch or split is in progress. If synchronizing, the status changes to PAIR at completion of the COPY; if splitting, the result is PSUS for primary disk, or SSUS for secondary disk.

Note If a backup is attempted while a disk is split (PSUS, SSUS), the backup fails with a status code 156. If a backup is attempted while a disk is in the COPY state, there are two possible results: if the disks synchronize (shown as PAIR), the backup proceeds; if the disks split (PSUS, SSUS), the backup fails with a status code 156.

%

Shows the percentage of the status that has completed.

P-LDEV#

The LDEV number of the “partner” disk in the pair.

M

Indicates whether the secondary is writable, as a result of being split from the primary.

▼ Configure the arrays, depending on your requirements

The next steps depend on the results of the `pairdisplay` listings and the requirements of your site.

- ◆ If all required disks are correctly paired (status of PAIR), the primary-secondary configuration is finished.
- ◆ If required disks are paired but currently split (PSUS, SSUS), or if they are not paired at all (SMPL), you must resynchronize or configure them, respectively.

Note If a mirror-type snapshot backup attempts to access a disk that is split or not paired, the backup fails with a status code 156.

- ◆ If disks are paired but need to be unpaired or otherwise reconfigured, you must split them and create a new association.



If disks are split:

1. Enter the following to resynchronize the split disks:

```
pairresync -g groupname -d dev_name
```

where *groupname* is the name listed under *dev_group*, and *dev_name* is the device name, as defined in the configuration files. To resynchronize the disks listed as split (PSUS, SSUS) in the above example (see “[Resulting output:](#)” on page 177), enter:

```
pairresync -g wiltest -d dev2
```

2. Enter the following to view the result:

```
pairdisplay -g wiltest -d dev2 -CLI -fc
```

When the resynchronization starts, the Status column reads COPY. When it is nearly completed, the Status column reads PAIR (see the % column for percentage completion).

If disks are not paired:

1. Enter the following to create a pair of primary and secondary:

```
paircreate -g groupname -d dev_name -vl
```

where *groupname* is the name listed under *dev_group*, *dev_name* is the device name, as defined in the configuration files, and -vl specifies that the current instance number is the primary.

To associate the dev3 disks as a pair (the ones listed as SMPL in the above example; see “[Resulting output:](#)” on page 177), enter the following:

```
paircreate -g wiltest -d dev3 -vl
```

2. Enter the following to view the result:

```
pairdisplay -g wiltest -d dev3 -CLI -fc
```

When the synchronization starts, the Status column reads COPY. When it is nearly completed, the Status column reads PAIR (see the % column for percentage completion).

If disks are paired but need to be split or reconfigured:

1. To split the secondary disk from the primary but maintain the pair association, enter the following:

```
pairsplit -g groupname -d dev_name
```

where *groupname* is the name listed under `dev_group`, and *dev_name* is the device name, as defined in the configuration files. The `pairdisplay` command will show a status of PSUS and SSUS.

For example:

```
pairsplit -g wiltest -d dev1
```

This splits the secondary from the primary in the `dev1` pair.

2. To split the secondary disk from the primary and remove the pair association between them, enter the following:

```
pairsplit -g groupname -d dev_name -S
```

where `-S` means break the pair association. The `pairdisplay` command will show SMPL in the P/S column for the affected disks, meaning the disks are no longer paired.

For more information on array configuration, refer to the documentation provided by the array's vendor.



Volume Manager Configuration

Disk Label

On Solaris only: If client data is configured in Volume Manager volumes, be sure to label all secondary devices using the `format` command (`label` option). Labeling the secondary disks prevents Volume Manager from marking the disks as disabled (if they are split from their primary disks) during a system reboot.

While a secondary disk is synchronized with its primary, the secondary is invisible to Volume Manager. When the secondary is split off from its primary disk, the secondary becomes visible again. If the secondaries are labeled (using the `format label` command), Volume Manager will not disable the disks when they are split.

Disk Types


There are important restrictions involving the use of Volume Manager with Advanced Client.

Note If these restrictions are not observed, the backup will fail.

- ◆ Do not include secondary (mirror) disks in a Volume Manager disk group.
- ◆ The Volume Manager disk group must contain disks of one vendor type only. Do not configure disks of different vendors in the same Volume Manager disk group.
- ◆ The vendor type of the snapshot method must match the vendor-type of the disks in the Volume Manager disk group.

Concerning these restrictions, refer to the next two diagrams.

Example VxVM Disk Groups: the Good and the Bad



P = primary disk in array

M = mirror (secondary)
disk in array

Consider the following VxVM disk groups:

If disk group contains P1, P2

Good: group contains only primary devices, of same vendor.

If disk group contains P3, M3

Bad: group contains a secondary (mirror) disk.

If disk group contains P1, P3, P5

Bad: group contains disks of different vendors.


As shown above, no secondary (mirror) disks should be included in VxVM disk groups, and groups must contain disks of the same vendor.

Note These restrictions apply when using any of the array-specific snapshot methods; they do NOT apply if you are using the **vxvm** snapshot method.



When Using Volume Manager and Array-Specific Methods

For each of these Volume Manager volumes:



select this snapshot method:



TimeFinder



ShadowImage



BusinessCopy

Disk Group Clones

When using the array-specific snapshot methods with client data configured over a Volume Manager volume, NetBackup creates a temporary disk group (clone) of the disks containing the mirror volume. To avoid a naming conflict in the Volume Manager, NetBackup names the temporary disk group as follows:

diskgroup_name_clone

For alternate client backup, the temporary disk group is named as follows:


client_name_diskgroup_name_clone

While the backup is in progress, this clone appears in the output of the Volume Manager `vxvg` command. This is normal. When the backup completes, NetBackup automatically removes the disk group clone.

Disk Group Cloning Example:

Client data is in:

- file system **/fs_1**
- configured over VxVM volume **/dev/vx/rdisk/dg_1/vol01**



VxVM disk group
dg_1 on
primary_disk1.



NetBackup creates a *temporary* VxVM disk group (clone) **client_name_dg_1_clone** on mirror_disk1.

In the above example, NetBackup removes the VxVM disk group *client_name_dg_1_clone* after the backup has completed. If a major system interruption occurs (such as a crash or unexpected reboot), NetBackup may not be able to remove the clone. In that case, you must use the `bpdgclone` command with the `-c` option to remove the clone, and then resynchronize the mirror disk with the primary disk. Refer to [“Removing a VxVM Volume Clone”](#) on page 226 for assistance.

When Secondary Disks are Split and Synched

This situation is described for your information only. A backup occurring in this circumstance should complete normally, in spite of the Volume Manager error described as follows.

When the secondary (mirror) device is split from its primary, Volume Manager will see the secondary disk as a separate device. But when the secondary disk is then re-synched to its primary disk (provided Volume Manager had seen it before), the synched secondary disk is no longer visible and VxVM issues an I/O error. In addition, if DMP is enabled, the secondary disks are marked as disabled. The next time the secondary is split, it will reappear in Volume Manager, only to disappear when the disk is again synched to its primary.



Best Practices

The recommendations in this section apply primarily to the use of the array-specific snapshot methods and Volume Manager, except where noted.

NetBackup Access to Arrays

In connection with the information listed under “[Access to Disk Arrays](#)” on page 165, note the following recommendation:

- ◆ The NetBackup media server only needs read access to the secondary disks in the array; it does not need access to the primary disks.

Resynchronizing Disks At End of Backup

Resynchronizing very large mirror disks can take time. If disk-resynchronization significantly delays completion of the backup, set the **Resynchronize mirror in background** option to Yes. This allows the backup to complete without waiting for the mirror disks to be resynchronized. The disks are resynchronized after the backup completes. Refer to [step 6](#) on page 89 for more information on this option.

Hardware-Level Disk Restore

Caution Hardware-level disk restore (such as by means of the `symmir` command with the `-restore` option) can result in data loss if the primary disk is shared by more than one file system or more than one VxVM volume. The hardware-level restore overwrites the entire primary disk with the contents of the mirror disk.

This can be a problem if you are attempting to restore a snapshot of *one* of the file systems or *one* of the VxVM volumes that share the same disk: the other file systems or volumes sharing the disk may have older data that you do not want to write back to the primary. When the hardware-level disk restore takes place, the older data will replace the newer data on the primary disk.

Volume Manager Disk Groups

When creating a VxVM disk group, it is best to create a group that corresponds to the primary disks that were grouped as described under “[Configuring Primary and Secondary Disks](#)” on page 167. If you create an array disk group with two primary disks, a VxVM disk group should be created with the same primaries. The VxVM disk group configuration should follow the array disk group configuration for the primaries.

Volume Manager with Dynamic Multipathing (DMP)

If you are using Volume Manager with DMP enabled, and there are multiple paths to the same disk array (for instance, one fibre channel connection and one SCSI), DMP will rename the array's disks with DMP encapsulated names.


Backups Concurrently Accessing Same Disk (no VxVM)

A conflict occurs if two or more backups using array-specific snapshot methods attempt to access a disk at the same time.

When the snapshot process is started, NetBackup reserves or “locks” the secondary (mirror) disk for that backup job, denying any other backup jobs access to that disk. If a second backup job requests a snapshot involving data on the same disk before the first job is complete, access is denied and the second job fails.

This conflict can arise when there are two backup policies, each using an array-specific method, and each requires access to the same disk at the same time (see diagram).

Backup Policies in Conflict: Two Backups Accessing Same Disk



Note Snapshot disk locks are applied to the entire disk: when a backup job requires a snapshot, the entire disk is locked.

To avoid this conflict, see “[Avoiding Concurrent Access Conflicts](#)” later in this chapter.



Backups Concurrently Accessing VxVM Volumes

A conflict occurs if two or more concurrent backups using an array-specific method attempt to access data configured in the same Volume Manager volume or in volumes configured on the same disk(s).

Concurrent Access to Same VxVM Volume

In this case, a conflict occurs if two or more backups using an array-specific method attempt to access the same Volume Manager volume at the same time.

Backup Policies in Conflict: Two Backups Accessing Same Volume



Backup policy A starts job to back up `/vol_1`; mirror is locked to make snapshot. Backup policy B attempts to back up `/vol_1` and requests snapshot: mirror is already locked—access denied.

The above diagram shows `/dev/vx/rdsk/dg/vol_1` on a single disk. The same conflict will occur if `/vol_1` is distributed across two or more disks.

To avoid this conflict, see “[Avoiding Concurrent Access Conflicts](#)” later in this chapter.


Concurrent Access to Volume Manager Volumes on Same Disks

A conflict can occur if two or more concurrent backups using an array-specific method attempt to access Volume Manager volumes that are distributed across the same disks.

VERITAS Volume Manager (VxVM) supports three means of distributing volumes across disks: striping, concatenating, and RAID 5 (described in the Volume Manager documentation). Use of these distribution methods can lead to access problems for NetBackup with an array-specific method. The following diagram shows two VxVM volumes, `/dev/vx/rdsk/dg/vol_1` and `/dev/vx/rdsk/dg/vol_2`. Each is distributed across two disks in an array (using any of the three distribution methods). If two backups request snapshots of these volumes at the same time, a conflict occurs, even

though the two backups are attempting to access different volumes. This happens because the array-specific snapshot methods split the mirror disk from the primary disk at the disk device layer, not at the volume layer.

Backup Policies in Conflict: Two Backups Accessing Volumes Distributed on Same Disks



Backup policy A starts to back up `/vol_1`; both disks A and B are locked to make a snapshot of `/vol_1`. Backup policy B attempts to back up `/vol_2` and requests snapshot: disks A and B are already locked—access denied.

Avoiding Concurrent Access Conflicts

These are recommendations for backups that encounter any of the concurrent-access problems when using an array-specific method.

- ◆ Schedule the policies so that none can start a backup at the same time as another.
- ◆ If possible, combine the separate policies into one policy. Snapshots will be created before backups begin, on a global basis across all streams at once.
- ◆ If you want the backups to run concurrently, combine the separate policies into one and configure that policy for multiple data streaming. Multiple data streaming prevents concurrent backups from encountering snapshot conflicts. See the *NetBackup System Administrator's Guide* for help with multiple data streams.
- ◆ If the data to back up is configured in Volume Manager (VxVM) volumes, use the vxvm snapshot method. The vxvm method allows snapshot backups to run concurrently without conflicts, *provided that the backup data consists of file systems mounted on VxVM volumes*. See [“Creating a Snapshot Mirror of the Source”](#) on page 150 for help with the vxvm method.
- ◆ Use the Volume Manager administration interface to determine which disks the volumes are configured on, and configure the volumes on different disks.





Notes on Media Server and Third-Party Copy

9

This chapter provides notes and restrictions regarding the NetBackup Media Server and Third-Party Copy Device methods.

The following topics are covered in this chapter:

- ◆ [Disk Requirements for Media Server/ Third-Party Copy](#)
- ◆ [ALL_LOCAL_DRIVES](#)
- ◆ [Storage Units](#)
- ◆ [Multiplexing](#)
- ◆ [Raw Partition Backups](#)



Disk Requirements for Media Server/ Third-Party Copy

For the NetBackup Media Server or Third-Party Copy Device backup method, the client's data must be on one or more disks that meet the following criteria:

- ◆ The disk must be either a SCSI or Fibre Channel device.
- ◆ The disk must be visible to both the NetBackup client and to the NetBackup media server. The disk must be connected through a fibre channel SAN or through a disk array that has dual port SCSI connections.
- ◆ The disk must be able to return its SCSI serial number in response to a serial-number inquiry (serialization), or the disk must support SCSI Inquiry Page Code 83.

ALL_LOCAL_DRIVES

The policy's Backup Selections list must not contain the ALL_LOCAL_DRIVES entry.

Storage Units

- ◆ Any_available is not supported for NetBackup Media Server and Third-Party Copy Device backup methods.
- ◆ Disk storage units are not supported for the Third-Party Copy Device method.

Multiplexing

The Third-Party Copy Device backup method is incompatible with multiplexing (the writing of two or more concurrent backup jobs to the same storage device). To prevent multiplexing on a third-party copy backup, you must set **Maximum multiplexing per drive** to 1 (on the "Add New Storage Unit" or "Change Storage Unit" dialog).

Raw Partition Backups

When entering a raw partition in the Backup Selections list (for a policy that uses the NetBackup Media Server or Third-Party Copy method), do not specify a *block device* as the raw partition. For these two backup methods, NetBackup does not support block devices. Instead, specify the raw partition as a *character device*.

Examples:

Solaris: /dev/rdisk/c1t3d0s3

HP: /dev/rdisk/c1t0d0

Backup and Restore Procedures

10

The following topics are covered in this chapter:

- ◆ [Performing a Backup](#)
- ◆ [Performing a Restore](#)
- ◆ [Configurations for Restore](#)
- ◆ [Restoring from a Disk Snapshot](#)



Performing a Backup

Before proceeding, please note the following for the array integration snapshot methods.

For the EMC **TimeFinder**, Hitachi **ShadowImage**, or HP **BusinessCopy** snapshot method, the client data to be backed up must reside on a mirror disk made by the corresponding vendor (EMC, Hitachi, or HP). Assistance from the disk array vendor's technical support may be required. For NetBackup-related items, refer to the chapter titled "[Snapshot Configuration Notes](#)."

Automatic Backup

The most convenient way to back up client data is to configure a policy and then set up schedules for automatic, unattended backups. To use NetBackup Advanced Client, you must enable snapshot backup as described in the appropriate configuration chapter of this guide. To add new schedules or change existing schedules for automatic backups, you can follow the guidelines in the *NetBackup System Administrator's Guide*.

Manual Backup

The administrator can use the NetBackup Administration interface on the master server to execute a backup for a policy. To use NetBackup Advanced Client, you must enable snapshot backup as described in the appropriate configuration chapter of this guide.

See the *NetBackup System Administrator's Guide* for instructions on doing manual backups.

User-Directed Backup and Archive

From a NetBackup client, the user can execute an Advanced Client backup. The NetBackup administrator must configure an appropriate snapshot policy with schedule.

See the *NetBackup Backup, Archive, and Restore Getting Started Guide* for instructions on doing user-directed backups and archives.

Performing a Restore

You can use the Backup, Archive, and Restore interface to restore individual files or directories, or a volume or raw partition. See the *NetBackup Backup, Archive, and Restore Getting Started Guide* for instructions on performing the restore. The following sections include restore notes and procedures unique to certain components of Advanced Client.

Restores from a FlashBackup Backup

Using the Backup, Archive, and Restore interface, you can restore individual directories or files (or an entire raw partition) from a FlashBackup backup. The procedure is the same as that for restoring from a regular backup as described in the *NetBackup Backup, Archive, and Restore Getting Started Guide*.

Note In the Backup, Archive, and Restore interface, set the policy type to **FlashBackup** for UNIX clients and **FlashBackup-Windows** for Windows clients.

Note the following before starting the restore.

- ◆ A backup made from a FlashBackup (or FlashBackup-Windows) policy supports both individual file restore and raw partition restore; that is, you can do either type of restore from the same backup. To restore individual files, select **Normal Backups** on the Restore Files tab; to restore an entire raw partition, select **Raw Partition Backups**.
- ◆ To restore a raw partition, you must have administrator capabilities on the NetBackup server.
- ◆ An entire raw partition can be restored from a full backup only. FlashBackup incremental backups only support individual file restores.
- ◆ Ensure that the device file for the raw partition exists prior to the restore.
- ◆ The overwrite option is ignored during raw partition restores. The device file must exist and the disk partition is overwritten whether this option is set or not.
- ◆ If you are restoring a very large number of files and individual file restore would take too long, you can do a raw-partition restore by redirecting it to another raw partition of the same size and then copying individual files to the original file system.
- ◆ For UNIX clients:
 - ◆ To restore an entire raw partition, ensure that the partition is *not* mounted and *not* in use. (For this reason, you cannot perform a raw partition restore to the root partition, or to the partition on which NetBackup is installed.) If the partition is being used by a database, shut down the database. The partition must be the same size as when it was backed up; otherwise, the results of the restore are unpredictable.



- ◆ After a raw partition restore of a VxFS file system, a file system consistency check (fsck) is usually required before the file system can be mounted.
- ◆ For Windows clients
 - ◆ To restore an entire raw partition, ensure that the partition *is* mounted (designated as a drive letter) but *not* in use. (For this reason, you cannot perform a raw partition restore to the root partition, or to the partition on which NetBackup is installed.) If the partition is being used by a database, shut down the database. The partition must be the same size as when it was backed up; otherwise, the results of the restore are unpredictable.

Restores in Clustered File System (VxFS on UNIX Only)

In a clustered file system, if you are restoring a large number of files (such as 100,000 or more), the restore will finish sooner if the file system is mounted on a local host. The file system can then be mounted as a shared mount after the restore.

Do the following:

1. Stop all applications (on any nodes) that are using the file system.
2. Unmount the file system.
3. Mount the file system locally.
4. Perform the restore.
5. Share the mounted file system again, and restart applications (if any).

Instant Recovery Restore

You can restore files from an Instant Recovery backup in the same way as from a normal backup. Restore procedures are described in the *NetBackup Backup, Archive, and Restore Getting Started Guide*.

In addition, there are several restore features unique to Instant Recovery.

- ◆ Block-level restore (for VxFS_Checkpoint snapshots)
- ◆ File promotion (for VxFS_Checkpoint or NAS_Snapshot snapshots)
- ◆ Fast File Resync for Windows (for vxvm and FlashSnap snapshots)
- ◆ Rollback (for VxFS_Checkpoint, vxvm, FlashSnap, or NAS_Snapshot snapshots)

These require special instructions, as follows.

Instant Recovery: Block-Level Restore (UNIX Clients Only)

If the Instant Recovery snapshot was made with the VxFS_Checkpoint method, large files can be recovered faster by means of block-level restore. Only the blocks that have changed are moved from the snapshot to the client's primary fileset.

Important Notes

- ◆ Block-level restore requires the VxFS File System.
- ◆ Block-level restore is available only when restoring files to the original location on the client, AND when the snapshot method used for the Instant Recovery backup was VxFS_Checkpoint.
- ◆ If the snapshot method for the backup was VxFS_Checkpoint and the files to be restored are in an Oracle database, block-level restore is automatically enabled.

To activate block-level restore

Create the following (empty) file on the client:

```
/usr/openv/netbackup/PFI_BLI_RESTORE
```

After this file is created, all subsequent restores of the client's data will use block-level restore.

To de-activate block-level restore

Delete (or rename) the PFI_BLI_RESTORE file.

Note When block-level restore is activated, it is used for all files in the restore. This may not be appropriate for all of the files. It may take longer to restore a large number of *small* files, because they must first be mapped.

Instant Recovery: File Promotion (UNIX Clients Only)

If the Instant Recovery snapshot was made with the VxFS_Checkpoint or NAS_Snapshot method for a UNIX client, large files that have had many changes since the backup can be recovered more quickly by means of file promotion. File promotion optimizes single-file restore by using a minimum of I/O to recover the files.

Notes on File Promotion

- ◆ Only regular files can be promoted, not file links or directories.
- ◆ **Notes on VxFS_Checkpoint:**



- ◆ File promotion requires the VxFS File System version 4.0 or later.
- ◆ File promotion can be done only from the last Instant Recovery snapshot that was made with the VxFS_Checkpoint method.
- ◆ File promotion is available only when restoring files to the original location on the original client.
- ◆ **Notes on NAS_Snapshot:**
 - ◆ File promotion is available when restoring to the original volume on the original client.
 - ◆ File promotion can be done from older snapshots, but any newer NAS snapshots are deleted after the file promotion takes place.
 - ◆ The file system requirements depend on the NAS vendor. For further requirements specific to your NAS vendor, see the *NetBackup for NDMP Supported OS and NAS Appliance Information* online document (refer to the preface for help accessing that document).

To Use File Promotion

The procedure for restoring individual files with file promotion is the same as the standard restore procedure for NetBackup (refer to the *NetBackup Backup, Archive, and Restore Getting Started Guide*). No special settings or choices are required when using the Backup, Archive, and Restore interface.

NetBackup Selects File Promotion on a File-by-File Basis

If the above requirements are met (see “[Notes on File Promotion](#)”), NetBackup automatically attempts file promotion for the file. Otherwise, the restore of the file takes place in the standard manner, without file promotion: all file data is copied from the snapshot to the primary file system. The NetBackup progress log indicates how many files were promoted, and how many files that could not be promoted were restored by means of tar.

Instant Recovery: Fast File Resync (Windows Clients Only)

If the Instant Recovery snapshot was made with the vxvm or FlashSnap method on a Windows client, large files that have had many changes since the backup can be recovered more quickly by means of Fast File Resync (a type of file promotion). Only the blocks that have changed are moved from the snapshot to the client’s primary fileset.

Notes on Fast File Resync (FFR)

- ◆ Requires Storage Foundations for Windows 4.1 or later and the licensed FlashSnap option.
- ◆ Can be done only from an Instant Recovery snapshot that was made with the vxvm or FlashSnap method.
- ◆ Available only when restoring to the original location on the original client.
- ◆ The “overwrite existing files” option must be selected.
- ◆ Regarding the files to be restored:
 - ◆ The original files must be present on the client, and will be overwritten by the restore.
 - ◆ The names and creation times of the original files and the snapshot files must be identical.
 - ◆ Files must be larger than 20 MB and be formatted in NTFS.
 - ◆ Files must not be compressed or encrypted.
 - ◆ There must be no open handles on either the original file or the snapshot file.

To Use Fast File Resync

The procedure for restoring individual files with fast file resync is the same as the standard restore procedure for NetBackup (refer to the *NetBackup Backup, Archive, and Restore Getting Started Guide*). No special settings or choices are required when using the Backup, Archive, and Restore interface.

NetBackup Selects Fast File Resync on a File-by-File Basis

If the above requirements are met (see “[Notes on Fast File Resync \(FFR\)](#)”), NetBackup automatically attempts fast file resync first. If fast file resync cannot be used for a given file, then the restore of that file takes place in the standard manner, without fast file resync: all file data is copied from the snapshot to the primary file system. The NetBackup progress log indicates how many files were resynced, and how many files that could not be resynced were restored by means of tar.

Instant Recovery: Snapshot Rollback

You can also restore a snapshot of an entire file system or volume with minimal I/O. This type of restore is called *point in time rollback*. All the data in the snapshot is restored; single file restore is not available in a rollback.



Notes on Rollback

- ◆ Rollback can be done only from backups that were enabled for Instant Recovery and made with the VxFS_Checkpoint, vxvm, FlashSnap, or NAS_Snapshot methods.
- ◆ For backups made with the VxFS_Checkpoint method, rollback requires the VxFS File System 4.0 or later and Disk Layout 6. For NAS_Snapshot, the file system requirements depend on the NAS vendor.
- ◆ Rollback deletes any VxFS_Checkpoint snapshots or NAS_Snapshot snapshots (and their catalog information) that were created *after* the creation-date of the snapshot that you are restoring.
- ◆ Rollback deletes all files that were created after the creation-date of the snapshot that you are restoring. Rollback returns a volume to a given point in time. Any data changes or snapshots that were made after that time are lost.
- ◆ Rollback is available only when restoring the file system or volume to the original location on the client.
- ◆ When a rollback of a file system is initiated, NetBackup verifies, by default, that the primary file system does not contain files created after the snapshot was made; otherwise, the rollback aborts.

▼ To perform snapshot rollback (UNIX)

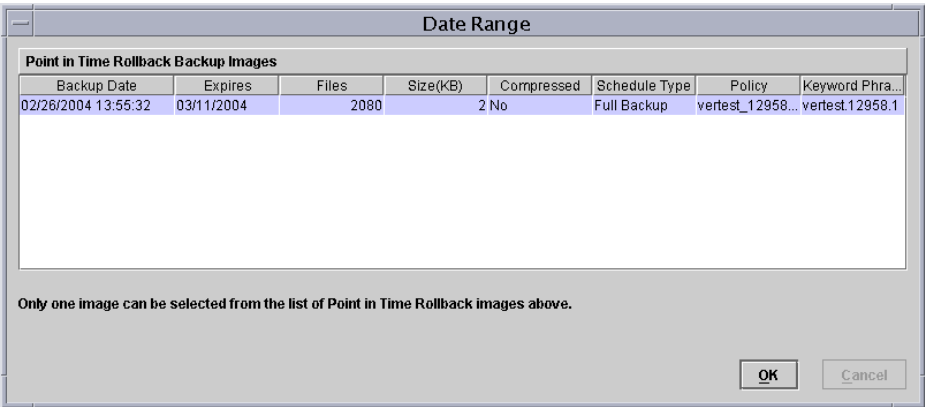
This procedure requires root access.

1. Start the Backup, Archive, and Restore interface.
`/usr/opensv/netbackup/bin/jbpSA &`
2. Click the **Restore Files** tab.
3. Click **Actions > Specify NetBackup Machines** to specify the server, source client, policy type, and destination client.
4. For the Restore Type, select **Point in Time Rollback**.

The **Browse directory** field is grayed out, with root (/) as default.



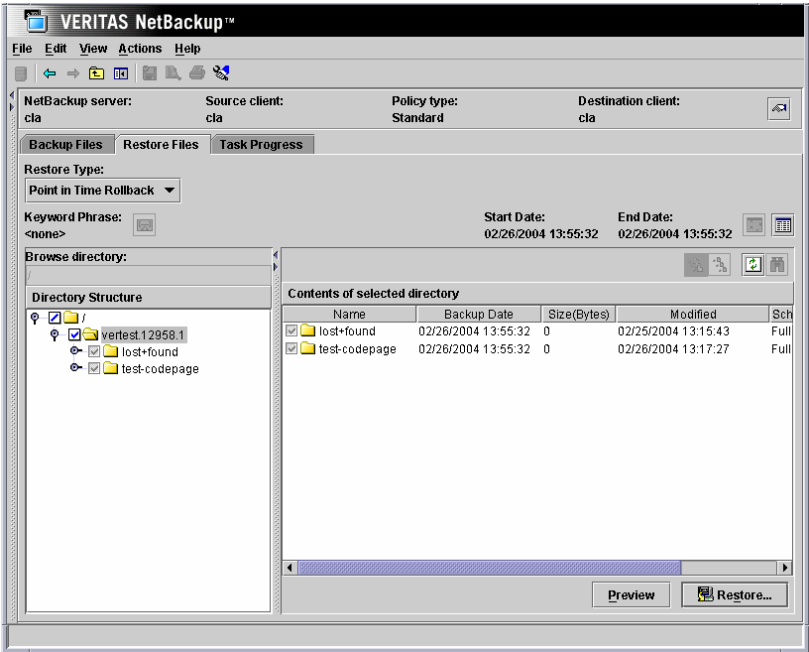
Instant Recovery backups are displayed in the Backup History window, for all dates (you cannot set a range).



- 5. Select an image from the list and click **OK**.

The image contents are displayed in the **Directory Structure** pane of the Restore Files tab.

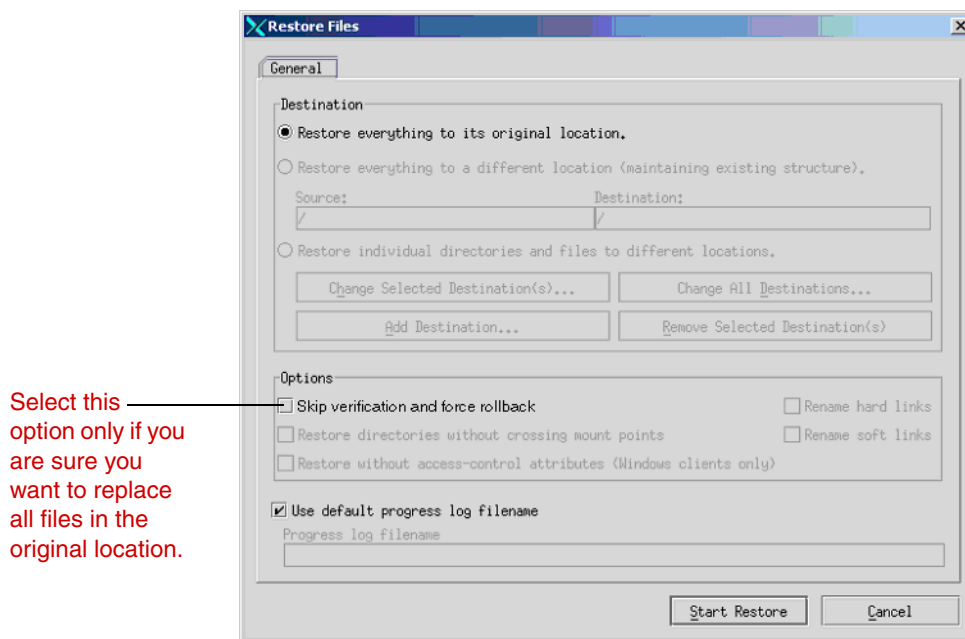
You can select root level or mount points (file systems or volumes), but not folders or files at a lower level



6. In the Directory Structure list, click the check box next to the root node or a mount point beneath root.

You can select a file system or volume, but not lower-level components.

7. Click the **Restore** button.



The only available destination option is **Restore everything to its original location**.

8. For file systems, you can choose to skip file verification by placing a check in the **Skip verification and force rollback** option.

Note For a rollback from a SnapVault, verification is always skipped, regardless of whether or not **Skip verification and force rollback** is selected; see “[Notes on SnapVault](#)” on page 121.

Caution Click on **Skip verification and force rollback** *only* if you are sure you want to replace all the files in the original location with the snapshot. Rollback deletes all files that were created after the creation-date of the snapshot that you are restoring.

If **Skip verification and force rollback** is not selected, NetBackup performs several checks on the file system as described under “[Notes on Rollback](#)” on page 200. If the checks do not pass, the rollback aborts and a message is written to the Task Progress tab stating that rollback could not be performed because file verification failed.

The rest of the procedure is identical to a normal restore as explained in the *NetBackup Backup, Archive, and Restore Getting Started Guide* and help.

▼ To perform snapshot rollback (Windows)

This procedure requires Administrator privilege.

1. Start the Backup, Archive, and Restore interface.

Click **Start > Programs > VERITAS NetBackup > Backup, Archive, and Restore**.

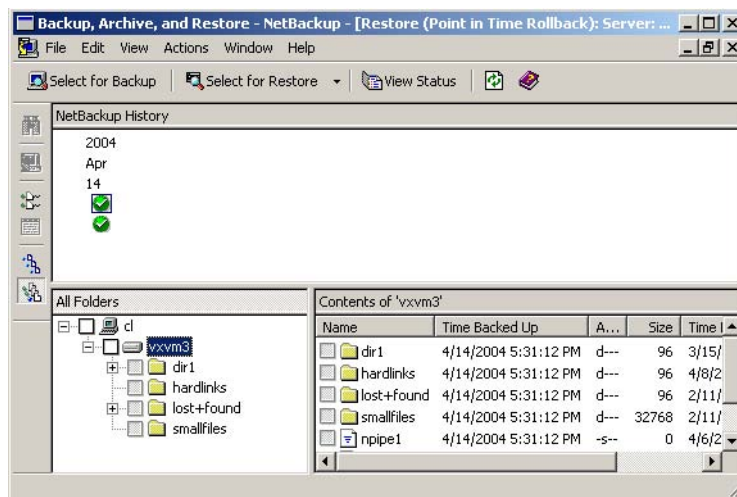
2. From the **Select for Restore** drop-down list, select **Restore from Point in Time Rollback**.

3. Click **File > Specify NetBackup Machines and Policy Type** to specify the server, source client, policy type, and destination client.

4. In the NetBackup History pane, click on the backup image to restore.

Only Instant Recovery backups are displayed in the NetBackup History pane, for all dates (you cannot set a date range).

You can select root level or mount points (file systems or volumes), but not folders or files at a lower level

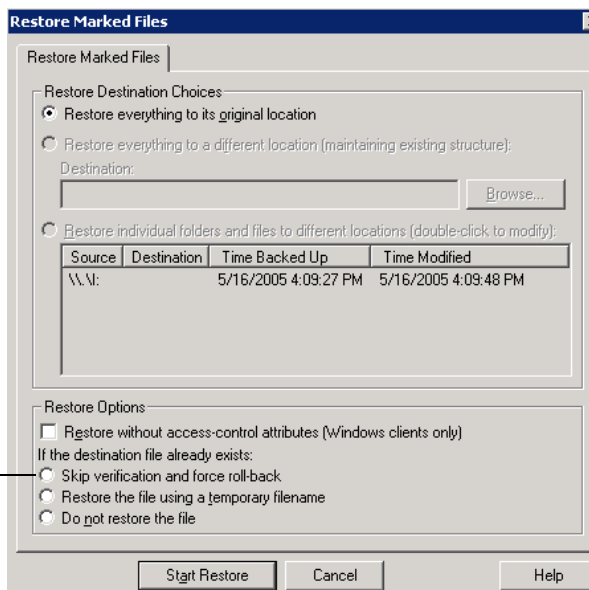


5. In the **All Folders** pane, click the check box next to the root node or a mount point beneath root.

You can select a file system or volume, but not lower-level components.

6. Click **Actions > Start Restore of Marked Files**.

Select this option only if you are sure you want to replace all files in the original location.



Note The only destination option is **Restore everything to its original location**.

7. For file systems, you can choose to skip file verification by placing a check in the **Skip verification and force rollback** option.

Note For a rollback from a SnapVault, verification is always skipped, regardless of whether or not **Skip verification and force rollback** is selected; see [“Notes on SnapVault”](#) on page 121.

Caution Click on **Skip verification and force rollback** *only* if you are sure you want to replace all the files in the original location with the snapshot. Rollback deletes all files that were created after the creation-date of the snapshot that you are restoring.

If **Skip verification and force rollback** is not selected, NetBackup performs several checks on the file system as described under “[Notes on Rollback](#)” on page 200. If the checks do not pass, the rollback aborts and a message is written to the progress log stating that rollback could not be performed because file verification failed.


The remainder of the procedure is identical to a normal restore as explained in the *NetBackup Backup, Archive, and Restore Getting Started Guide*.



Configurations for Restore


Advanced Client backups can be restored in several ways, depending on your configuration:

- ◆ Restore over the LAN: data can be restored from the storage device to the media server, and from the media server via the LAN to the client. This kind of restore is also used for ordinary (non-Advanced Client) backups.



1. Media server reads data from local storage.
2. Media server sends the data to the client over the LAN.
3. Client restores the data to disk (disk can be locally attached or on SAN).

- ◆ Restore over the SAN to a host acting as both client and media server: this requires the `FORCE_RESTORE_MEDIA_SERVER` option in the server's `bp.conf` file (see the *NetBackup System Administrator's Guide* for details on this option).



1. Client/media server reads data from tape over the SAN.
2. Client restores the data to disk (disk can be locally attached or on SAN)
(Requires use of `FORCE_RESTORE_MEDIA_SERVER` option in `bp.conf` file.)

- ◆ Restore directly from a snapshot (this is not the Instant Recovery feature): if the “Keep snapshot after backup” option was turned on for the backup, the data can be restored from a mirror disk by restoring individual files from the snapshot, or by restoring the entire snapshot. Note that this type of restore must be done from the command prompt (using, for instance, a copy command such as UNIX `cp`), not from the NetBackup Administration Console.

For details, refer to “[Restoring from a Disk Snapshot](#)” on page 207.

Restoring from a Disk Snapshot

If the “Keep snapshot after backup” parameter is set to Yes (on the Snapshot Options dialog), the snapshot is retained on the mirror disk after the backup completes. From the command line, you can restore individual files or the entire snapshot directly from the disk, rather than restoring from tape.

Note Unless the backup was made with the Instant Recovery feature, you cannot restore from a snapshot by means of the Backup, Archive, and Restore interface; you must perform the restore manually at the command line.

On UNIX

▼ To restore individual files

To restore individual files, locate and mount the snapshot file system and copy files from that file system using system commands, such as `cp` and `ftp`, as follows.

1. To list the identifiers of current snapshots, use the `bpfis query` command with the `query` option:

```
/usr/opensv/netbackup/bin/bpfis query
```

This returns the ID (FIS IDs) of all current snapshots. For example:

```
INF - BACKUP START 3629
INF - FIS IDs: 1036458302
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

2. For each snapshot identifier, enter `bpfis query` again, specifying the snapshot ID:

```
/usr/opensv/netbackup/bin/bpfis query -id 1036458302
```

This returns the path of the original file system (snapshot source) and the path of the snapshot file system. For example:



```
INF - BACKUP START 3629
INF - Snapshot image host : ricopico
INF - Snapshot image owner: GENERIC
INF - Time created       : Mon Oct  7 19:35:38 2002
INF - REMAP FILE BACKUP /mnt/ufscon USING /tmp/_vrts_frzn_img_26808/mnt/ufscon
OPTIONS:ALT_PATH_PREFIX=/tmp/_vrts_frzn_img_26808,FITYPE=MIRROR,
MNTPOINT=/mnt/ufscon,FSTYPE=ufs
INF - EXIT STATUS 0: the requested operation was successfully completed
```

In this example, the primary file system is `/mnt/ufscon` and the snapshot file system is `/tmp/_vrts_frzn_img_26808/mnt/ufscon`.

3. Copy the files from the mounted snapshot file system to the original file system.

▼ To restore the entire snapshot

There are additional ways to recover data from the disk snapshot, depending on your hardware configuration and the snapshot method used by the policy.

If the snapshot method was FlashSnap, you can restore the snapshot volume as follows:

1. Unmount the snapshot source (original file system) and the snapshot file system on the alternate client:

```
umount original_file_system
umount snapshot_image_file_system
```

To locate the file systems, refer to [step 1](#) and [step 2](#) on page 207.

2. Deport the snapshot on the alternate-client:

- a. Find the VxVM disk group:

```
vxdbg list
```

The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

If `vxdbg list` does not show the disk group, the group might have been deported. You can discover all the disk groups, including deported ones, by entering:

```
vxdisk -o alldgs list
```

The disk groups listed in parentheses are not imported on the local system.

- b. Deport the VxVM disk group:

```
vxdbg deport SPLIT-primaryhost_diskgroup
```


3. Import and join the VxVM disk group on the primary (original) client:

```
vxdbg import SPLIT-primaryhost_diskgroup
vxrecover -g SPLIT-primaryhost_diskgroup -m
vxdbg join SPLIT-primaryhost_diskgroup diskgroup
```

**4. Start the volume and snap back the snapshot volume, using the
-o resyncfromreplica option:**

```
vxvol -g SPLIT-primaryhost_diskgroup start SNAP_diskgroup_volume
vxassist -g SPLIT-primaryhost_diskgroup -o resyncfromreplica snapback SNAP_diskgroup_volume
```

If the snapshot was made on an EMC, Hitachi, or HP disk array:

WITH CAUTION, you can use hardware-level restore to restore the entire mirror or secondary disk to the primary disk. If the disk is shared by more than one file system or VxVM volume, there may be unintended results. *Please read the following:*

Caution Hardware-level disk restore (such as by means of the `symmir` command with the `-restore` option) can result in data loss if the primary disk is shared by more than one file system or more than one VxVM volume. The hardware-level restore overwrites the entire primary disk with the contents of the mirror disk.

This can be a problem if you are attempting to restore a snapshot of *one* of the file systems or *one* of the VxVM volumes that share the same disk: the other file systems or volumes sharing the disk may have older data that you do not want to write back to the primary. When the hardware-level disk restore takes place, the older data will replace the newer data on the primary disk.

On Windows

▼ **To restore individual files**

1. To list the identifiers of current snapshots, use the `bpfis` command with the `query` option:

```
/usr/opensv/netbackup/bin/bpfis query
```

This returns the ID (FIS IDs) of all current snapshots. For example:

```
INF - BACKUP START 3629
INF - FIS IDs: 1036458302
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

2. For each snapshot identifier, enter `bpfis query` again, specifying the snapshot ID:



```
/usr/opensv/netbackup/bin/bpfiis query -id 1036458302
```

This returns the path or the original file system (snapshot source) and the GUID (Global Universal Identifier) representing the snapshot volume. For example:

```
INF - BACKUP START 2708
INF - Snapshot method: FlashSnap
INF - Snapshot image host : tire
INF - Snapshot image owner : NBU
INF - Time created       : Sat Oct 25 15:26:04 2003

INF - REMAP FILE BACKUP H:\ USING
\\?\Volume{54aa666f-0547-11d8-b023-00065bde58d1}\
OPTIONS:ALT_PATH_PREFIX=C:\Program Files\VERITAS\NetBackup\
Temp\_vrts_frzn_img_2408,FITYPE=MIRROR,MNTPOINT=H:\,FSTYPE=NTFS
INF - EXIT STATUS 0: the requested operation was successfully completed
```

In this example the snapshot file system is H:\ and the GUID is \\?\Volume{54aa666f-0547-11d8-b023-00065bde58d1}\.

3. To restore individual files from the snapshot volume:

a. Mount the GUID to an empty NTFS directory:

```
mountvol C:\Temp\Mount
\\?\Volume{54aa666f-0547-11d8-b023-00065bde58d1}\
```

b. Copy the file to be restored from the temporary snapshot mountpoint (in this example, C:\Temp\Mount) to the primary volume.

▼ To restore the entire snapshot

There are additional ways to recover data from the disk snapshot, depending on your hardware configuration and the snapshot method used by the policy.

If the snapshot method was FlashSnap, you can restore the snapshot volume as follows:

1. Deport the snapshot on the alternate-client:

a. Find the VxVM disk group:

```
vxldg list
```

The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

b. Deport the VxVM disk group:

```
vxldg -g split_diskgroup deport
```

2. Import and join the VxVM disk group on the primary (original) client:

```
vxassist rescan  
vxdg -g split_diskgroup import  
vxdg -g split_diskgroup -n diskgroup join
```

3. Snap back the snapshot volume, using the `-o resyncfromreplica` option:

```
vxassist -o resyncfromreplica snapback \Device\HarddiskDmVolumes\diskgroup\snap_volume
```



This chapter covers the following topics.

- ◆ [Gathering Information and Checking Logs](#)
- ◆ [Latest Patches and Updates](#)
- ◆ [Important Notes](#)
- ◆ [Particular Issues](#)
- ◆ [Removing a Snapshot](#)
- ◆ [Removing a VxVM Volume Clone](#)

Note For explanations of NetBackup status codes, refer to the “NetBackup Status Codes and Messages” chapter in the *NetBackup Troubleshooting Guide*.



Gathering Information and Checking Logs

You can resolve many problems on your own by creating the logging directories listed below, reproducing the problem, and then checking the logs. For an in-depth description of NetBackup logs, refer to the *NetBackup Troubleshooting Guide*.

Logging Directories for UNIX Platforms

Advanced Client messages are written to the following directories if they exist (see table below).

Note Messages pertaining to NAS_Snapshot and SnapVault can be found in the ndmp unified log (originator ID 151), in `/usr/opensv/logs`.

Note To create detailed log information, place a VERBOSE entry in the `bp.conf` file on the NetBackup master and client, or set the Global logging level to a high value in the Logging dialog, under both Master Server Properties and Client Properties.

Note These directories can eventually require a lot of disk space. Delete them when you are finished troubleshooting and remove the VERBOSE option from the `bp.conf` file or reset the Global logging level to a lower value.

UNIX Logging Directories for Backup

During a backup, Advanced Client messages are logged to the following directories. Create these directories using an access mode of 755 so NetBackup can write to the logs. You can use the `/usr/opensv/netbackup/logs/mklogdir` script to create these directories.

Path of log directory	Where to create the directory
<code>/usr/opensv/netbackup/logs/bpbrm</code>	NetBackup master server if Instant Recovery backup to disk only; otherwise, on media server
<code>/usr/opensv/netbackup/logs/bptm</code>	NetBackup media server
<code>/usr/opensv/netbackup/logs/bpbkar</code>	NetBackup client or alternate client
<code>/usr/opensv/netbackup/logs/bpfis</code>	NetBackup client or alternate client

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bppfi	NetBackup client or alternate client

UNIX Logging Directories for Restore

During a restore, Advanced Client messages are logged to the following directories on the master server. Create these directories using an access mode of 755.

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bprestore	NetBackup master server
/usr/opensv/netbackup/logs/bprd	NetBackup master server
/usr/opensv/netbackup/logs/bpbrm	NetBackup master server if Instant Recovery backup to disk only; otherwise, on media server
/usr/opensv/netbackup/logs/bptm	NetBackup media server

snaptcl Driver Messages

Messages from the `snaptcl` driver are logged in the client's `/var/adm/messages` file along with other kernel messages.

Logging Folders for Windows platforms

During a backup, Advanced Client messages are written to the folders listed below if they exist. You can use the following command to create these folders:

```
install_path\NetBackup\logs\mklogdir.bat
```

The default path for the logs is `C:\Program Files\VERITAS\NetBackup\logs`. Since a different path can be set during installation, the path listed below is shown as *install_path*\NetBackup\logs.

Note To create detailed log information, set the Global logging level to a high value, in the Logging dialog, under both Master Server Properties and Client Properties.



Note The log folders can eventually require a lot of disk space. Delete them when you are finished troubleshooting and set the logging level on master and client to a lower value.

Windows Logging Folders for Backup

During a backup, Advanced Client messages are logged to the following folders.

Note Messages pertaining to NAS_Snapshot and SnapVault can be found in the ndmp unified log (originator ID 151), in *install_path*\NetBackup\logs.

Path of log directory	Where folder is created
<i>install_path</i> \NetBackup\logs\bpbm	NetBackup master server if Instant Recovery backup to disk only; otherwise, on media server
<i>install_path</i> \NetBackup\logs\bptm	NetBackup media server
<i>install_path</i> \NetBackup\logs\bpffi	NetBackup client or alternate client
<i>install_path</i> \NetBackup\logs\bpbkar	NetBackup client or alternate client
<i>install_path</i> \NetBackup\logs\bpfis	NetBackup client or alternate client

Windows Logging Folders for Restore

During a restore, Advanced Client messages are logged to the following folders on the master server.

Path of log directory	Where folder is created
<i>install_path</i> \NetBackup\logs\bprestore	NetBackup master server
<i>install_path</i> \NetBackup\logs\bpri	NetBackup master server

Path of log directory	Where folder is created
<i>install_path</i> \NetBackup\logs\bpbrm	NetBackup master server if Instant Recovery backup to disk only; otherwise, on media server
<i>install_path</i> \NetBackup\logs\bptm	NetBackup media server

Contacting VERITAS Customer Support

Before calling customer support, please gather as much log information as possible. Be sure to have the following information ready:

- ◆ NetBackup version
- ◆ Operating system version of the NetBackup master and media server and NetBackup Advanced Client client
- ◆ Note whether or not the action that failed had ever worked, and whether the problem is repeatable
- ◆ Log information

Latest Patches and Updates

When using other VERITAS products such as the VERITAS File System and Volume Manager, or Storage Foundation, it is a good idea to install the latest patches and updates for those products. Installing the latest software can fix a variety of issues.

For example:

If you receive status code 156 (snapshot error encountered) when running a FlashSnap alternate client backup, and the client data is configured in Storage Foundations for Windows 4.1, the volume to be backed up may appear to be missing on the alternate client even though the disk group has been split on the primary client and the disk deported to the alternate client. Installing the latest Storage Foundation patches has been known to correct this problem.



Important Notes

- ◆ If backup or restore jobs are running slowly, verify that the network interface cards (NIC) are set to full duplex. Half duplex often causes poor performance. For assistance viewing and resetting duplex mode for a particular host or device, consult the documentation provided by the manufacturer. You may be able to use the `ifconfig` (or `ipconfig`) command to view and reset duplex mode, as explained in the *NetBackup Troubleshooting Guide*.
- ◆ The disk containing the client's data (the files to back up) must be a SCSI or Fibre Channel device if you are using NetBackup Media Server or Third-Party Copy Device.
- ◆ The disk containing the client's data must be visible to both the client and the media server if you are using the NetBackup Media Server or Third-Party Copy Device method. The disk can be connected through SCSI or fibre channel.
- ◆ For the NetBackup Media Server or Third-Party Copy Device method, a disk device must be able to return its SCSI serial number in response to a serial-number inquiry (serialization), or the disk must support SCSI Inquiry Page Code 83.
- ◆ When configuring the Third-Party Copy Device or NetBackup Media Server method, a particular storage unit or group of storage units must be specified for the policy—do not choose **Any_available**. For configuration instructions, refer to [“Configuring an Advanced Client Policy”](#) on page 79.
- ◆ The `storage_unit_name` portion of a `mover.conf.storage_unit_name` file name must exactly match the actual storage unit name (such as `nut-4mm-robot-t14-0`) that you have defined for the policy. See [“Naming the Mover File”](#) on page 68 for help creating a `mover.conf.storage_unit_name` file.

Similarly, the `policy_name` portion of a `mover.conf.policy_name` file name must match the actual name of the policy that the third-party copy device is to be associated with.
- ◆ For the **TimeFinder**, **ShadowImage**, or **BusinessCopy** snapshot methods, the client data must reside in a device group, with the data on the primary disk and synchronized with a mirror disk. Assistance from the disk array vendor may also be required. Refer to [“Array-Related Snapshot Methods”](#) on page 161.
- ◆ If the “Keep snapshot after backup” option for the snapshot method is changed from yes to no, the last snapshot created for that policy must be deleted manually before the backup is run again. Use the `bpfis` command to delete the snapshot. Refer to the man page for `bpfis`.
- ◆ During a third-party copy device backup, if tape performance is slow, increase the buffer size by creating one of the following files on the media server:

```
/usr/opensv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC.policy_name  
/usr/opensv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC.storage_unit_name
```

```
/usr/opensv/netbackup/db/config/SIZE_DATA_BUFFERS_TPC
```

By default, the size of the data buffer for third-party copy backup is 65536 bytes (64K). To increase it, put a larger integer in the `SIZE_DATA_BUFFERS_TPC` file. For a buffer size of 96K, put 98304 in the file. If the value is not an exact multiple of 1024, the value read from the file will be rounded up to a multiple of 1024.

The file name with no extension (`SIZE_DATA_BUFFERS_TPC`) applies as a default to all third-party copy backups, if neither of the other file-name types exists. A `SIZE_DATA_BUFFERS_TPC` file with the `.policy_name` extension applies to backups executed by the named policy, and the `.storage_unit_name` extension applies to backups using the named storage unit. If more than one of these files applies to a given backup, the buffers value is selected in this order:

`SIZE_DATA_BUFFERS_TPC.policy_name`

`SIZE_DATA_BUFFERS_TPC.storage_unit_name`

`SIZE_DATA_BUFFERS_TPC`

As soon as one of the above files is located, its value is used. A `.policy_name` file that matches the name of the executed policy will override the value in both the `.storage_unit_name` file and the file with no extension. The `.storage_unit_name` file will override the value in the file with no extension.

You can set the maximum buffer size that a particular third-party copy device can support.

Note A third-party copy device will not be used if it cannot handle the buffer size set for the backup.

Particular Issues

Installation Problem

If you receive the following message during installation:

```
/usr/opensv/netbackup/bin/version not found.  
Add-On Product Installation Aborted.
```

you have tried to install the Advanced Client software before installing the base NetBackup software.



FlashBackup and Status code 13

Status 13 can result from any of the following:

- ◆ The FlashBackup cache partition may have run out of space. In that case, the cache partition may not be large enough for the backup requirements.

If the cache partition is full, you will see messages such as the following in the system log:

```
WARNING: sn_alloccache: cache /dev/vx/rdsk/flashldg/f full - all  
snaps using this cache are now unusable  
WARNING: sn_failsnap: snapshot id 5 failed error 28
```

Specify a larger cache partition, or designate additional cache partitions in the Backup Selections list. See the “[FlashBackup Configuration](#)” chapter in this guide for cache partition requirements.

- ◆ **On Solaris:** If your cache partition runs out of space, there may be “stale” snapshots taking up space on the cache partition. Stale snapshots are those that were not automatically deleted by FlashBackup.
 - a. Determine if there are stale snapshots on your Solaris client by executing the following:

```
/usr/opensv/netbackup/bin/driver/snaplist
```
 - b. For each snapshot listed, execute the following to make sure there is a bpbkar process associated with it:

```
ps -eaf |grep ident
```

where *ident* is the snapshot process id displayed by the snaplist command.
 - c. Remove snapshots that do not have an associated bpbkar process by entering the following:

```
/usr/opensv/netbackup/bin/driver/snapoff snapn
```

where *snapn* is the snapshot id displayed by the snaplist command.

Removing a Snapshot

NetBackup ordinarily removes snapshots after the Advanced Client backup completes, unless the “Keep snapshot after backup” parameter was set to Yes. However, as a result of certain kinds of system failures, such as a system crash or abnormal termination of the backup, the snapshot may not be removed.

▼ To identify and remove a left-over snapshot

1. Use the `bpfis` command with the `query` option to list the current snapshots:

Do this on the client or alternate client, depending on the type of backup:

```
/usr/opensv/netbackup/bin/bpfis query
```

This returns the IDs (FIS IDs) of all current snapshots. For example:

```
INF - BACKUP START 3629
INF - FIS IDs: 1036458302
INF - EXIT STATUS 0: the requested operation was successfully
completed
```

In this example, the snapshot ID is 1036458302.

2. If the `bpfis` output shows the ID of the snapshot, delete it as follows:

```
bpfis delete -id snapshot_id
```

If `bpfis` removed the snapshot, you can skip the rest of this procedure.

3. Solaris, HP, AIX, Linux: if `bpfis` could not remove the snapshot, enter the following (on the client or alternate client) when no backups are running:

```
df -k
```

This displays all mounted file systems, including any snapshots of a mounted file system.

Note It is important to enter `df -k` when no backups are running. If a snapshot backup is currently running, the snapshot should not be deleted. NetBackup will delete it when the backup completes.

Here are two snapshots from a `df -k` listing:

```
/dev/dsk/clt3d2s4 1048800 73076 914742 8% /tmp/_vrts_frzn_img__wil_vxfs_1299000
/dev/vx/dsk/clone_ges_clone/ufs 38383 21678 12867 63% /tmp/_vrts_frzn_img__mix_ufs_1299000
```

The snapshot appears in the following form:

```
/tmp/_vrts_frzn_img__filesystemname_pid
```



4. Solaris, HP, AIX, Linux: unmount the unneeded snapshot file systems (on the client or alternate client, depending on the type of backup).
5. The next step depends on the type of snapshot.

For nbu_snap (Solaris only):

- a. Enter the following to display leftover snaps:

```
/usr/opensv/netbackup/bin/driver/snaplist
```

- b. To remove a leftover snap, enter

```
/usr/opensv/netbackup/bin/driver/snapoff snap1 ... snapn
```

For more information on `naplist` and `napoff`, refer to the “[Managing nbu_snap \(Solaris\)](#)” appendix of this guide.

For vxvm (Solaris, HP, AIX, Linux) and VVR (Solaris and HP):

Do the following on the client for vxvm, and on the alternate client for VVR:

- a. Enter the following to display unsynchronized mirror disks:

```
vxprint -g diskgroup
```

- b. Enter the following to resynchronize the mirror disks:

```
vxassist -g diskgroup -v volume snapback
```

For vxvm (Windows):

- a. Enter the following to display unsynchronized mirror disks:

```
vxdg -g diskgroup dginfo
```

- b. Enter the following to resynchronize the mirror disks:

```
vxassist snapback  
\\Device\\HarddiskDmVolumes\\diskgroup\\snap_volume
```

For VxFS_Checkpoint (Solaris, HP, AIX, Linux):

- a. Enter the following VxFS command to display the name of the checkpoint:

```
/usr/lib/fs/vxfs/fsckptadm list /file_system
```

where *file_system* is the mount point of the primary or original file system that was backed up, NOT the snapshot file system that was unmounted in [step 4](#).

For example, if the snapshot file system that was unmounted is the following:

```
/tmp/_vrts_frzn_img_vm2_1765
```

The original file system, which should be specified on the `fsckptadm list` command, is this:

```
/vm2
```

Example entry:

```
/usr/lib/fs/vxfs/fsckptadm list /vm2
```

Output:

```
/vm2
NBU+2004.04.02.10h53m22s:
  ctime           = Fri Apr 02 10:53:23 2004
  mtime           = Fri Apr 02 10:53:23 2004
  flags           = removable
```

In this example, the name of the checkpoint is NBU+2004.04.02.10h53m22s.

- b.** Remove the checkpoint by entering the following:

```
/usr/lib/fs/vxfs/fsckptadm remove name_of_checkpoint /file_system
```

For example:

```
/usr/lib/fs/vxfs/fsckptadm remove NBU+2004.04.02.10h53m22s /vm2
```

- c.** If the checkpoint cannot be removed, unmount it (`umount`) and retry [step b](#).

For more detail on removing VxFS clones, refer to the recommended actions for NetBackup status code 11 in the *NetBackup Troubleshooting Guide*.

For TimeFinder, ShadowImage, BusinessCopy (Solaris or HP only):

Do the following on the client or alternate client, depending on the type of backup:

- a.** To discover and remove any VxVM clones, follow the steps under “[Removing a VxVM Volume Clone](#)” on page 226.
- b.** Enter the following to resynchronize the mirror disks:

For EMC arrays (TimeFinder):

```
symmir -g device_group establish LdevName
```

where *LdevName* is the logical device name of the standard device.



For Hitachi and HP arrays (ShadowImage, BusinessCopy):

```
pairresync -g groupname -d dev_name
```

For more information about EMC, Hitachi, and HP arrays and resynchronizing disks, refer to the section “[Array-Related Snapshot Methods](#)” on page 161.

For VxFS_Snapshot (Solaris or HP only):

- ◆ Using the mounted file system found at [step 3](#), unmount the snapshot as follows:

```
umount -F vxfs /tmp/_vrts_frzn_img__filesystemname_pid
```

For FlashSnap (Solaris, HP, AIX, Linux):

Do the following on the client or alternate client, depending on the type of backup:

- a. Find the VxVM disk group:

```
vxvg list
```

The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

If `vxvg list` does not show the disk group, the group might have been deported. You can discover all the disk groups, including deported ones, by entering:

```
vxdisk -o alldgs list
```

The disk groups listed in parentheses are not imported on the local system.

- b. Deport the VxVM disk group:

```
vxvg deport SPLIT-primaryhost_diskgroup
```

Enter the following on the primary (original) client:

- c. Import and join the VxVM disk group:

```
vxvg import SPLIT-primaryhost_diskgroup  
vxrecover -g SPLIT-primaryhost_diskgroup -m  
vxvg join SPLIT-primaryhost_diskgroup diskgroup
```

- d. Start the volume and snap back the snapshot volume:

```
vxvol -g SPLIT-primaryhost_diskgroup start SNAP-diskgroup_volume  
vxassist snapback SNAP-diskgroup_volume
```


Example

In this example, `chime` is the primary client and `rico` is the alternate client. `lhddg` is the name of the original disk group on `chime`, and `chime_lhddg` is the split group that was imported on `rico` and must be rejoined to the original group on the primary `chime`.

On alternate client `rico`, enter:

```
vx dg deport chime_lhddg
```

On primary client `chime`, enter:

```
vx dg import chime_lhddg
vxrecover -g chime_lhddg -m
vx dg join chime_lhddg lhddg
vxvol start SNAP-lhddg-vol01
vxassist snapback SNAP-lhddg-vol01
```

For FlashSnap (Windows):

- a. Find the VxVM disk group:

```
vx dg list
```

The format of the disk group name is as follows:

```
SPLIT-primaryhost_diskgroup
```

- b. Deport the VxVM disk group:

```
vx dg -g split_diskgroup deport
```

Enter the following on the primary (original) client:

- c. Import and join the VxVM disk group:

```
vxassist rescan
vx dg -g split_diskgroup import
vx dg -g split_diskgroup -n diskgroup join
```

- d. Snap back the snapshot volume:

```
vxassist snapback \Device\HarddiskDmVolumes\diskgroup\snap_volume
```



Removing a VxVM Volume Clone

A form of snapshot that might need manual deletion is a VxVM volume clone. See “[Disk Group Clones](#)” on page 184 for a description of disk clones.

Major system interruptions, such as a system crash or unexpected reboot, may prevent NetBackup from removing the clone. If the clone is not removed, subsequent backups of the client’s data will fail. Examine the `/usr/opensv/netbackup/logs/bpfis` log for text such as the following:

```
19:13:07.686 [14981] <2> onlfi_vfms_logf: INF - do_cmd: Command failed with status=20:
/usr/opensv/netbackup/bin/bpdgclone -g wil_test -n vol01 -f /var/tmp/HDSTFCAAs7aOqD
</dev/null >/var/tmp/VfMSAAQ7aOqD 2>/var/tmp/VfMSBAAr7aOqD
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF - --- Dumping file /var/tmp/VfMSAAQ7aOqD
(stdout):
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF - --- End of file /var/tmp/VfMSAAQ7aOqD
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF - --- Dumping file /var/tmp/VfMSBAAr7aOqD
(stderr):
19:13:07.687 [14981] <2> onlfi_vfms_logf: INF -      clone group and volume already exists
19:13:07.688 [14981] <2> onlfi_vfms_logf: INF - --- End of file /var/tmp/VfMSBAAr7aOqD
```

In this case, you must use the `bpdgclone` command with the `-c` option to remove the clone, and then resynchronize the mirror disk with the primary disk.

▼ How to remove the clone

The following commands should be run on the client or alternate client, depending on the type of backup.

1. When no backups are running, use the following VxVM command to list any clones.

```
vxpdg list
```

Note If a backup configured with an array-specific snapshot method is currently running, a clone for that backup will appear in the `vxpdg` output. Do not delete the clone; NetBackup will delete it when the backup completes.

Example `vxpdg` output:

NAME	STATE	ID
rootdg	enabled	983299491.1025.turnip
VolMgr	enabled	995995264.8366.turnip
wil_test_clone	enabled	1010532924.21462.turnip
wil_test	enabled	983815798.1417.turnip

In this example, the name suffix indicates `wil_test_clone` was created for a snapshot backup that was configured with an array-specific snapshot method. If a backup failed with log entries similar to those included above, the clone must be manually deleted.

2. To remove the clone, enter the following:

```
/usr/opensv/netbackup/bin/bpdgclone -g disk_group -n volume -c clone
```

For the above example, you would enter:

```
/usr/opensv/netbackup/bin/bpdgclone -g wil_test -n vol01 -c wil_test_clone
```

where `wil_test` is the name of the disk group, `vol01` is the name of the VxVM volume, and `wil_test_clone` is the name of the clone. Use the Volume Manager `vxprint` command to display volume names and other volume information.

For more information, refer to the `bpdgclone` man page. For assistance with `vxprint` and other Volume Manager commands, refer to the *VERITAS Volume Manager Administrator's Guide*.

3. To verify that the clone has been removed, re-enter `vx dg list`.

Sample output:

NAME	STATE	ID
rootdg	enabled	983299491.1025.turnip
VolMgr	enabled	995995264.8366.turnip
wil_test	enabled	983815798.1417.turnip

The clone no longer appears in the list.



Processing Background

A

This appendix presents background information on snapshot processing and copy-on-write snapshots.

Note For a diagram of the daemons/services involved in creating and backing up a snapshot, refer to the Functional Overview appendix of the *NetBackup Troubleshooting Guide*.

- ◆ [Processing Before and After the Snapshot](#)
- ◆ [How Copy-on-Write Works](#)



Processing Before and After the Snapshot

NetBackup performs several vital functions prior to creating a snapshot, as outlined below and in the following text. Without this pre-processing, the integrity of the snapshot cannot be guaranteed and the backup data may be of no value.

NetBackup Processing Before and After Creating the Snapshot

Steps 1, 2, and 6 apply only to databases, such as those requiring NetBackup for Oracle Advanced Client.

1. Backup process requests database *quiesce*.
2. Database application quiesces (must wait for transactions to complete).
3. Lock and flush the file system.
4. Create the snapshot.
5. Unlock the file system.
6. Release (unquiesce) the application.
7. Back up the snapshot.
8. (Optional:) Remove the snapshot.

Quiescing the System

Before a useful snapshot can be created, the data to back up must be transactionally consistent or complete. A transaction is a single data action, such as updating a patient's record in a medical database, or creating a record for a new patient. Such a transaction is composed of multiple I/O requests (search, copy, send, write, and so forth). Until the transaction's I/O requests are complete, the data is inconsistent and may be unsuitable for backup.

Transactions affect all levels of the storage management stack (file system, volume manager, and so forth), generating further transactions as a request is handed off to the next level of the stack. From the viewpoint of the file system, for instance, an I/O request issued by a database application constitutes a transaction and may be split into many disk references, all of which must be complete for the original request to be fulfilled. As a result, the creation of the snapshot must be coordinated with any application or other process that can affect the transactional consistency of the data.

The means of coordination is called *quiesce* (literally, to make quiet or place in repose). This involves pausing the database application or process until the data is transactionally consistent. Applications and the storage management stack must all be quiesced before a useful snapshot can be made.

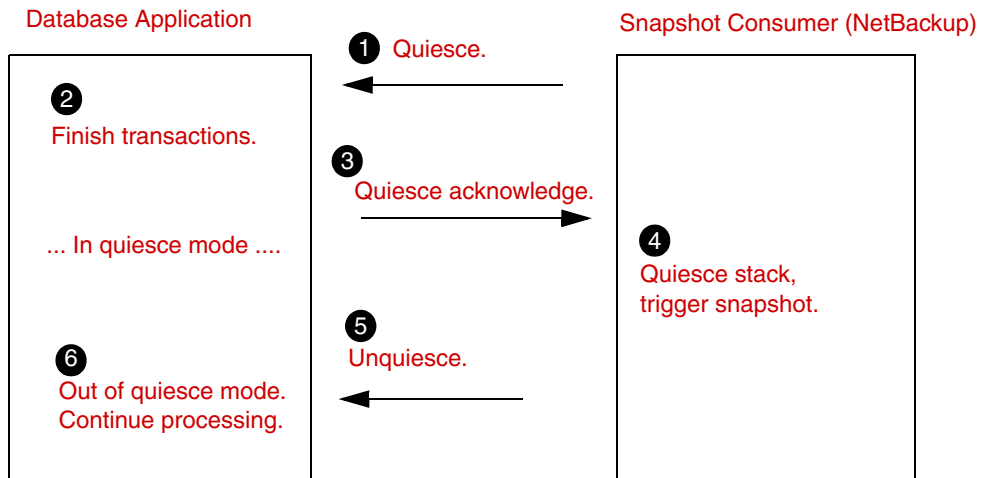
Quiescing the Database Application

Most database applications are transactionally consistent only at particular points in time. Sometimes, they are consistent only after they have been shut down. Since there is a growing need for database applications to remain up and available constantly, many applications are now designed to reach a point of transactional consistency at regular intervals or in response to an external event. This process is called *application quiesce*, described below.

In database application quiesce, an external signal or message is sent to a receptive database. In response, the database finishes the current transaction or group of transactions, signaling the snapshot consumer when this is complete. The database then waits for a second signal indicating that normal operations can resume. After the database signals that it has reached a state of transactional consistency, the final steps of creating the snapshot can proceed.

Once the snapshot has been created, another signal is sent to the waiting database telling it to resume processing. This is called *unquiescing* the application.

Dialog for Quiesce/Unquiesce



Quiescing the Stack

The storage management stack is a layered arrangement of software elements. An I/O request originated by a database application passes from element to element until a hardware request to move data reaches the storage network. Each stack element performs a variety of functions, some of which treat I/O requests like transactions to assure their completion. Before a snapshot is created, the stack must be quiesced (made transactionally consistent).

Since the file system is the front-line interface to applications for managing files and performing I/O, file system quiesce is a critical part of quiescing the stack.

File System

Two of the principal tasks of quiescing the file system are the following:

- ◆ Prohibit new I/O requests from initiating, which is called *locking the file system*.
- ◆ *Flush* file system cache (write cached data back to disk). The system must complete any outstanding application I/O and note completion of outstanding metadata updates.

Volume Manager

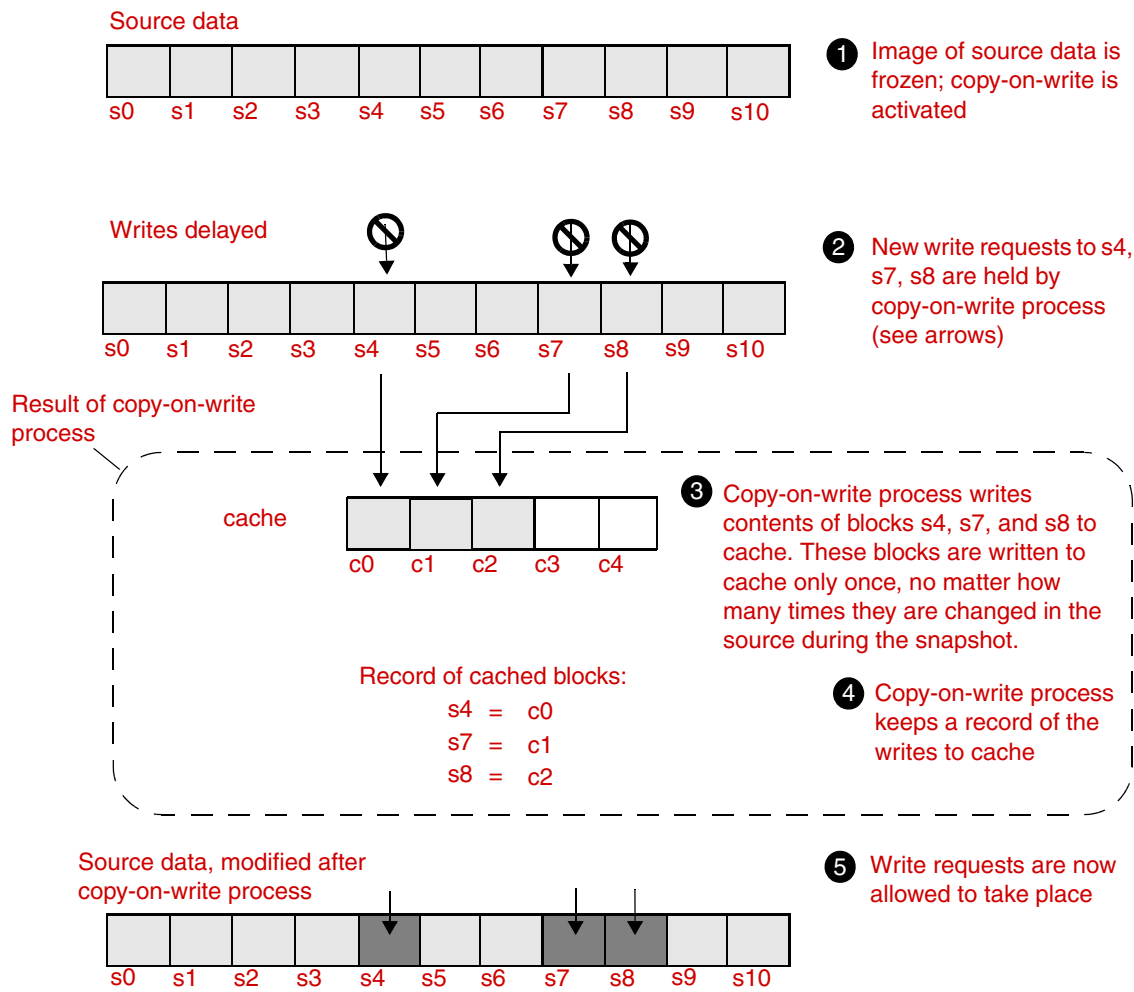
As in a file system, the volume manager's data caching may have to be flushed and disabled until the snapshot is created. As long as volume manager caching is enabled, data required for a consistent image may be lingering in volume manager cache rather than being available on disk when the snapshot is created.

How Copy-on-Write Works

A copy-on-write type of snapshot is a detailed account of data as it existed at a certain moment. Unlike a mirror, a copy-on-write is not really a *copy* of the data, but a specialized "record" of it.

The copy-on-write process works as follows: when a snapshot is required, any unfinished transactions or changes to the source data are allowed to complete, but new changes are temporarily stalled. The source is momentarily idled (made quiescent). Once the copy-on-write is activated, new transactions or changes (writes) to the source data are allowed to take place. However, the copy-on-write process briefly intercepts or holds the first write request that is issued for any particular block of data. While holding those requests, it copies to cache the blocks that will be affected by those writes, and keeps a record of the cached blocks. In other words, it reads each source block that is about to change for the first time, copies the block's current data to cache, and records the location and identity of the cached blocks. Then the intercepted writes are allowed to take place in the source blocks. (See figure "[Copy-on-write process](#)" on page 233.)

Copy-on-write process

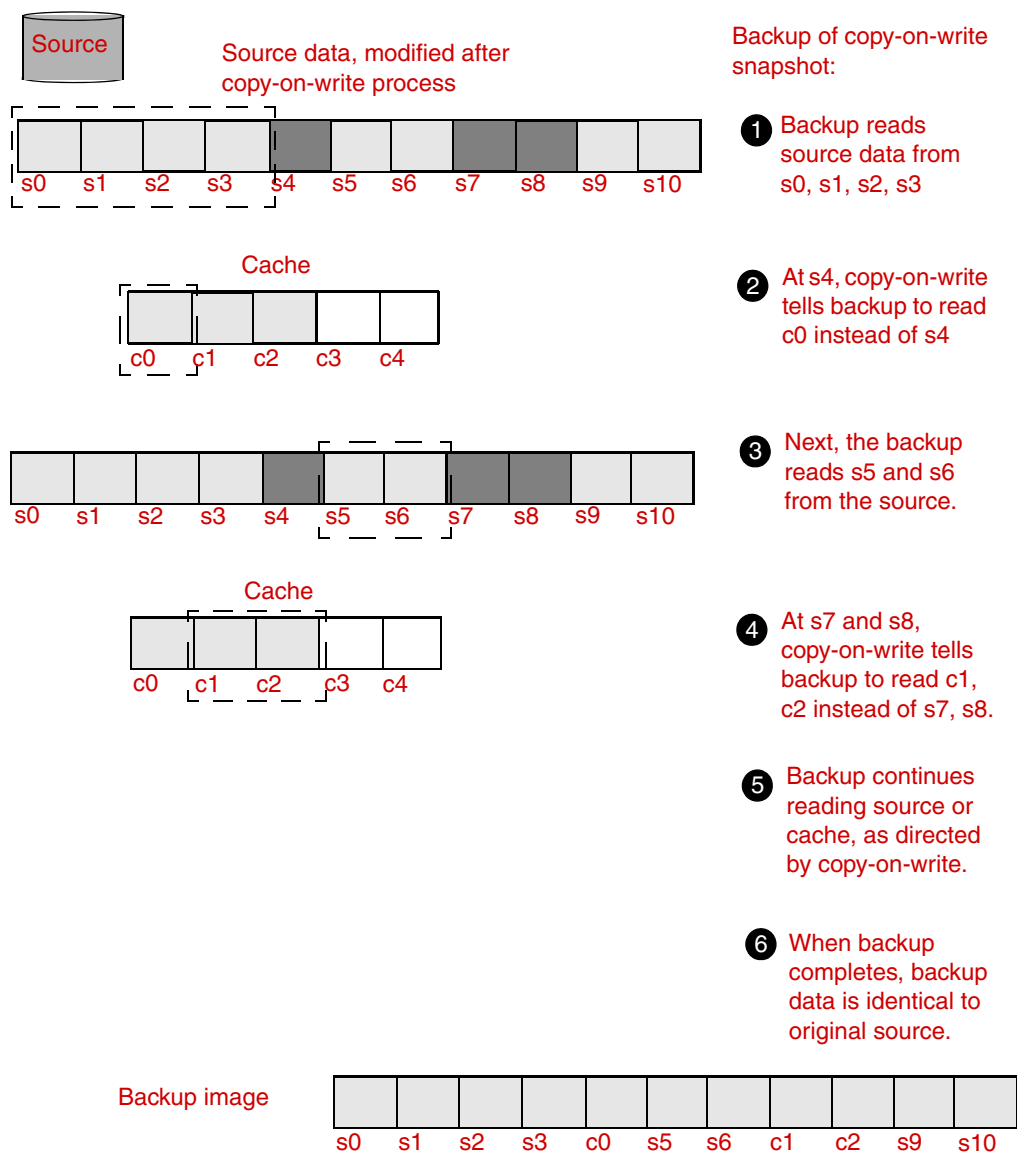


The immediate results of the copy-on-write are the following: a cached copy of those portions of the source that were about to change at a certain moment (see step 3 above), and a record of where those cached portions (blocks) are stored (4).

The copy-on-write does not produce a copy of the source; it creates cached copies of the blocks that have changed and a record of their location. The backup process refers to the source data or cached data as directed by the copy-on-write process (see next diagram).



Backing up a copy-on-write snapshot



As shown in the above diagram, an accurate backup image is obtained by combining the unchanged portions of the data with the cache. When a backup of the snapshot begins, the backup application copies the source data **1** until it comes to a block that changed after the copy-on-write process was activated. The copy-on-write tells the backup to skip that

changed block and read in its place the cached (original) copy ❷. The backup application continues copying source data ❸ until it comes to another changed block. Cache is read again ❹ as the copy-on-write process dictates. The backup, when finished, is an exact copy of the source as it existed the moment the copy-on-write was activated.



Managing nbu_snap (Solaris)

This appendix describes commands for managing snapshots that were made with the `nbu_snap` method.

Note This appendix applies to Solaris systems only.



Cache for nbu_snap

The `nbu_snap` method is used by NetBackup to create a copy-on-write snapshot in the following circumstances (Solaris systems only):

- ◆ If the **nbu_snap** method is configured for the back up on the Snapshot Options dialog.
- ◆ If the **nbu_snap** method was chosen by NetBackup's auto-selection mechanism when the backup was initiated.
- ◆ If the client is in a FlashBackup policy and no snapshot method is configured (**Perform snapshot backups** on the Policy display is not selected). This is called FlashBackup default mode, and will be discontinued in a future release.

In all cases, a raw partition must be specified as cache for the snapshot. The same cache can be used by any number of concurrent `nbu_snap` backups, as long as the cache is large enough to hold copies of all file system blocks that were changed by user activity while the snapshots were active.

In the Administration Console, the raw partition cache can be specified in any of three places (see "[How to Enter the Cache](#)" on page 145).

Determining Cache Size

The size needed for the cache partition depends on the amount of user write activity that occurs while the snapshot is active, not on the size of the file system. The more user activity, the larger the cache must be. You can use the `snplist` and `snpcachelist` commands described in this chapter to determine the size of the cache required for that activity. For a procedure, see "[Sizing the Cache Partition](#)" on page 143.

If a cache is too small and overflows, all snapshots using the cache become unreadable and the backups that are reading the snapshots will fail.

Terminating nbu_snap

NetBackup ordinarily removes snapshots after the Advanced Client backup completes. However, as a result of certain kinds of system failures, such as a system crash or abnormal termination of the backup, the snapshot may not have been removed.

Use the `snapoff` command to terminate an `nbu_snap` snapshot that was not terminated by the backup job (see "[snapoff](#)" on page 242). For more information on terminating snapshots, see "[Removing a Snapshot](#)" on page 221.

nbu_snap Commands

The following commands relate to the nbu_snap snapshot method.

snapon

snapon starts an nbu_snap snapshot (copy-on-write).

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapon snapshot_source cache
```

where *snapshot_source* is the partition on which the client's file system (the file system to be "snapped") is mounted, and *cache* is the raw partition to be used as copy-on-write cache.

Example 1:

```
/usr/opensv/netbackup/bin/driver/snapon /var /dev/rdisk/c2t0d0s3
```

Example 2:

```
/usr/opensv/netbackup/bin/driver/snapon /dev/vx/rdsk/omo/tcp1
/dev/vx/rdsk/omo/sncache
```

Note The snapshot is created on disk, and remains active until it is removed with the snapoff command or the system is rebooted.

snaplist

snaplist shows the amount of client write activity that occurred during an nbu_snap snapshot. Information is displayed for all snapshots that are currently active.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snaplist
```

Information is displayed in the following form:

id	ident	size	cached	minblk	err	time
9	6730	2097152	2560	0	0	05/29/03 07:52:18
	device = /dev/vx/rdsk/omaha/tcp1					
	cache = /dev/vx/rdsk/omaha/sncache					
16	7857	2097152	128	0	0	05/29/03 12:16:00
	device = /dev/vx/rdsk/omaha/vol01					
	cache = /dev/vx/rdsk/omaha/sncache					
17	7908	20971520	4224	0	0	05/29/03 12:17:38
	device = /dev/vx/rdsk/omaha/vol03					
	cache = /dev/vx/rdsk/zetab/cache					



If no snapshots are currently active, no output is returned.

Description:

◆ `id`

The snapshot ID. This is used by the `snappoff` command to terminate the snapshot.

◆ `ident`

A unique numeric identifier of the snapshot (this is the pid of the process that created the snapshot).

◆ `size`

The size of the client's snapshot source in 512-byte blocks. This is the partition on which the client's file system (the file system being backed up) is mounted. Note: this size value is not a reliable guide as to the size of cache needed for the snapshot; the user write activity during the snapshot is what determines the size of the cache needed. See the `cached` column of this output.

◆ `cached`

The number of 512-byte blocks in the client file system that were changed by user activity while the snapshot was active. Prior to being changed, these blocks were copied to the cache partition. The more blocks cached as a result of user activity, the larger the cache partition required. However, additional overhead—not shown in this `cached` value—is required in the cache.

Note To see the total space used in a particular cache partition, use the `snappcachelist` command.

◆ `minblk`

In the partition on which the file system is mounted, `minblk` shows the lowest numbered block that is currently being monitored for write activity while the snapshot is active. `minblk` is used by FlashBackup policies only.

◆ `err`

An error code; 0 indicates no error.

If a snapshot has encountered an error, the `err` will be non-zero and the snapshot will be inaccessible. It can be terminated using `snappoff` and the snapshot ID. Error codes are identified in `/usr/include/sys/errno.h`. Also, error messages may be found in `/var/adm/messages`.

◆ `time`

The time at which the snapshot was started.

◆ `device`

The raw partition containing the client's file system data to back up (snapshot source).

- ◆ `cache`

The raw partition used as cache by the copy-on-write snapshot process.

Note Make sure this partition is large enough to store all the blocks likely to be changed by user activity during the backup. To determine the total space used in a particular cache partition by all active snapshots, use the `snapcachelist` command.

snapcachelist

`snapcachelist` displays information about all partitions currently in use as `nbu_snap` caches. This command shows the extent to which the caches are full.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapcachelist
```

If no snapshots are currently active, no output is returned.

Output is of the form:

device	free	busy
/dev/rdsd/c0t4d0s0	238528	264472

Description:

- ◆ `device`

The raw partition being used as cache.

- ◆ `free`

The number of 512-byte blocks unused in the cache partition.

- ◆ `busy`

The number of 512-byte blocks in the client data that changed while the snapshot was active. Prior to being changed, these blocks were copied to the cache partition by the `nbu_snap` copy-on-write process. For each cache device listed, `busy` shows the total space that was used in the cache. You can use this value as a sizing guide when setting up raw partitions for `nbu_snap` cache.

When a cache is full, any additional change to the client data will cause the copy-on-write to fail and the snapshot will no longer be readable or writable. Reads or writes to the client data will continue (that is, user activity will be unaffected). The failed snapshot, however, will not be terminated automatically and must be terminated using `snapoff`.



Note `snaplist` and `snapcachelist` can also be used to monitor an `nbu_snap` snapshot that was started by a NetBackup policy. Once the backup completes, NetBackup removes the snapshot. As a result, the `snaplist` and `snapcachelist` commands no longer return any output.

snapstat

`snapstat` displays diagnostic information about the snap driver.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapstat
```

snapoff

`snapoff` terminates an `nbu_snap` snapshot.

Execute this command as root:

```
/usr/opensv/netbackup/bin/driver/snapoff snap1 ... snapn
```

where *snappx* is the numeric id of each copy-on-write process to be terminated. (Use `snaplist` to display the id of active snapshots.)

If `snapoff` is successful, a message of the following form will be displayed:

```
snapshot 1 disabled
snapshot 2 disabled
...
snapshot n disabled
```

If `snapoff` fails, an explanatory message is displayed. Error codes are identified in `/usr/include/sys/errno.h`.

Caution Do not terminate a snapshot while the backup is active, because corruption of the backup image may result.

Index

Numerics

- 156 status code 217
- 3pc.conf file 54, 69
 - creating 69
 - description 54
 - x option 42

A

- a=wwpn
 - lun 56
- abnormal termination 221, 238
- access conflicts (arrays) 187, 188, 189
- access time not updated 98
- accessibility features xv
- activate block-level restore 197
- Active Directory 100, 160
 - and FlashBackup 100
- active third-party copy device 41
- actual device file name 103
- address (in 3pc.conf) 56
- address mapping mode 47
- Administrator account and NAS_Snapshot 113
- Advanced Client 4
 - access web info xvi
 - installing 30
 - products included in 2
- Advanced Snapshot Options button 84
- AIX
 - and VxFS 29
 - media servers, restrictions 22
- ALL_LOCAL_DRIVES entry 21, 81, 97, 103, 135, 192
- Allow multiple data streams 97, 189
- alternate client
 - defined 23
 - notes on 93
- alternate client backup 3, 9, 102
 - and disk arrays 166

- and FlashSnap 153
- and split mirror 11
- and VVR 156
- configuring 93
- disk config requirements 165
- disk group naming 184
- introduction 10
- requirements 93
- restrictions 93, 125
- testing setup 153, 157, 159
- alternate world-wide name 57
- Any_available storage unit 81, 97, 218
- APP_NBU_VVR 156
- arbitrated loop 24
- archive bit
 - incremental backup 94
- archives 194
- Array Integration Option (see Advanced Client)
- arrays
 - access conflicts 187
 - and VxVM 182
 - as third-party copy devices 74
 - create config files 171
 - list of supported 163
 - multi-ported 18
 - restrictions 182
 - view state of 177
- associating primary to mirror 167
- auto snapshot method
 - and Instant Recovery 132
- auto snapshot selection 84, 87, 94, 132
- automatic backup 194

B

- background on processing 229
- backup
 - agent 2, 23
 - automatic 194



- expiration 129
- frequency 130
- local 8
- logs 214, 216
- manual 194
- of copy-on-write 234
- offhost
 - configuration 80
 - prerequisites 192
 - SAN configuration for 37
- raw partition 82, 192
- retention period
 - NAS_Snapshot 158
- scripts 92
- SCSI offhost 17
- techniques (overview) 4
- types supported 102
- user-directed 194
- Backup Policy Configuration wizard 82, 97, 135
- Backup Selections list 81
 - ALL_LOCAL_DRIVES entry 21, 81, 97, 103, 135, 192
 - and Instant Recovery 134
 - block vs. character device 192
 - directives 106
 - FlashBackup 103
 - symbolic link 96
- Backup, Archive, Restore 200
- Bare Metal Restore 80
- BCV 164
- best practices 186
- BLIB 78
- block device file (vs character) 192
- block level incremental backup 78
- block-level restore 196, 197
 - how to activate 197
 - restriction 197
- bp.conf file 214
- bpbkar
 - log 214
 - process 220
- bpbrm log 214, 215
- bpdgclone command 227
- bpfis
 - command 207, 209, 221
 - log 214
- bpmoverinfo command 72
- bpffi log on client 215

- bprd log 215
- bprdreq -terminate 31
- bprestore log 215
- bpSALInfo
 - address info in 3pc.conf 56
- bptm log 63, 214, 215
- bptpcinfo command 42, 43, 49, 69, 172
- BusinessCopy method 87, 162, 194, 218

C

- c=client 56
- cache 89
 - definition of 23
 - diagram of 233, 234
 - flushing for quiesce 232
 - object 136, 150
 - overflow 238
 - partition 106
 - partition out of space 220
 - requirements 142
 - size 238
 - specifying raw partition for 142
- CACHE= directive 104, 106, 146
- cfigmgr 48
- Change Policy dialog 86, 92
- Chaparral 62
- character device file (vs block) 142, 192
- checking logs 214
- checkpoint
 - removing (VxFS) 222
 - VxFS file system 147
- client data
 - prerequisites for offhost backup 192
- client list 103
- client name (in 3pc.conf) 56
- client option (on bptpcinfo) 42, 69
- client path (in 3pc.conf) 56
- client software
 - distributing 31
- clients
 - installing software 31
- clone
 - removing 222, 226
 - VxVM disk group 184, 226
- clustering 31
- CommandCentral 39, 49, 50, 81
- Common Serial Number Mode 167
- concurrent access conflicts (arrays) 187, 189
- configuration



- auto method 84
- backup selections 103
- client list 103
- files 72
- FlashBackup 101
- prerequisites 78
- procedure 79, 86, 131
- robot in Media Manager 113
- snapshot wizard xvi, 79, 113, 131, 135, 136, 137, 145
- supported data types 100
- copy manager (see third-party copy)
- copy-on-write
 - cache 233
 - compared to mirror 5
 - defined 23
 - for Storage Checkpoint 127
 - how backed up 234
 - overview 5, 232
- Core Frozen Image Services (see Advanced Client)
- cp command 207
- cross mount points (disabled) 96
- Crossroads 62
- customer support 217

D

- data change object 136
- data consistency (quiesce) 230
- data mover (see third-party copy)
- Database Edition for Oracle 21
- database files (Windows) 100, 160
- data-property-name entry 51
- DB2 124
- DCO plexes 153
- deinstallation
 - clients 36
 - server 35
- deleting
 - a clone 226
 - snapshots 128
- deport
 - disk group 153, 154
 - snapshot 208, 210
- device
 - ID 56
 - path 56
 - recognition 48
 - serialization 22, 192, 218

- Device Configuration wizard 78
- device file name (FlashBackup) 103
- device path, passthru 65
- directives
 - for scripts 92
- directives (Backup Selections list) 106
- directories for troubleshooting 214
- disk
 - access conflicts 187
 - restoring snapshot from 207
 - SCSI vs. IDE 218
 - visibility 218
- disk array
 - as third party copy device 74
 - choose snapshot method 162
 - create config files 171
 - supported 163
- disk group 182
 - clone 184, 226
 - shared in VxVM 142, 153
 - VxVM 186
- DISK keyword 64
 - passthru driver device path needed 65
- disk layout for Storage Checkpoint 125
- disk restore
 - caution 186, 209
- disk snapshot
 - restoring from 207
- disk storage units 81, 192
- distributing software to clients 31
- DLL patch for VxFS 21
- DMP (Dynamic Multipathing) 167, 185, 187
- DMX disk groups
 - creating 168
- dr keyword 67
- duplex mode and performance 218

E

- E4 target (see identification descriptor)
- EMC CLARiON 163
- EMC Symmetrix/DMX disk groups
 - creating 168
- END keyword 66
- exceptions to exclude list 100
- exclude list 100, 103
- expiring backup images 129
- Extended Copy command 3, 9, 16, 27, 54
- Extended Frozen Image Services (see Advanced Client)



extent, defined 24

F

failover 31
Fast File Resync (in Storage Foundation for
 Windows 4.1) 133, 196, 198
fast mirror resynch 151
FastResync 24, 26, 133, 153
feature/requirements table 19
FFR (Fast File Resync) 133, 196, 198
fibre channel
 defined 24
 types supported 24
file pathname (max length) 82, 96
file promotion 4, 196, 197, 198
file systems
 defined 24
 quiescing 232
files list (see Backup Selections list)
FIXED address mapping mode 47
FlashBackup 2, 238
 actual device name 103
 and VxFS_Checkpoint 147
 backup selections 103
 device file name 103
 features 100
 files to back up 103
 restoring 195
FlashSnap method 21, 87, 133, 199
 and status 156 217
 deporting the disk 153
 preparing for alternate client backup 153
 restoring from disk image 208, 210
 with VxVM shared disk group 153
flush file system 230, 232
FORCE_RESTORE_MEDIA_SERVER
 option 206
format command 182
fsck (after raw partition restore of vxfs
 system) 196
fsckptadm command 27, 127, 222
full-sized instant snapshots 89, 137, 151, 152

G

get_license_key 30
GID 93
glossary of terms 23
GROUP keyword 65
group of storage units 81

H

hardware-level disk restore 186, 209
HBA drivers 53
Hitachi 62, 87
HOMRCF 87
HORCM_CMD 172
HORCM_DEV 173
HORCM_INST 173
HORCM_MON 171
HORCMINST 176
horcmstart command 176
HP 62, 106
 8 devices per target 48
hr keyword 64, 67

I

I/O components
 supported 100
i=iddesc 58
i=reserve value 66
IBC messages 156
IBC send/receive timeout (seconds) 91
IDE vs. SCSI 218
identification descriptor 58, 81
importing disk group 154, 155
include list 103
incremental backup
 archive bit and timestamp 94
inetd 175
initbprd 32
Inline Tape Copies (Vault) 22
inraid command 172
inquiry page code 83 22, 192, 218
insf command 48
install script 30
install_adc 36
installation
 clients 31
 deinstallation 35
 directories 30
 mixed environments 34
 of Advanced Client 30
 upgrade (patch) 31
 upgrading 35
Installation Options menu 30
instance number 176
Instant Recovery 4
 defined 24
 deleting snapshots 128



- Fast File Resync 133, 196, 198
 - Schedule tab 134
 - selecting in policy 132
 - snapshot only 132
 - volume name restriction 125, 135, 150
 - instant snapshots 89, 136, 137, 150, 151
 - ioscan command 43, 48
- J**
- JBOD array 48
 - jbpSA 200
 - jnbSA 79, 86, 92, 101, 131
- K**
- Keep snapshot after backup 90, 221
 - restoring from image 207
 - kernel messages 215
 - keywords in mover.conf 64, 65
- L**
- l=lun 57
 - label option of format command 182
 - labeling secondary disks 182
 - left over snapshot 221, 238
 - license keys, installing 30
 - Limit jobs per policy 97
 - limitations (also see the Release Notes) 21
 - links (in Backup Selections list) 96
 - Linux
 - and VxFS 29
 - local host backup method
 - network configuration for 8
 - local system account 113
 - lock file system 230, 232
 - logging
 - directories to create 214
 - logical volume (as raw partition) 142
 - logs 214, 215, 216
 - creating for UNIX 214
 - creating for Windows 215
 - for SCSI reserve/release 63
 - loop (Fibre Channel) 24
 - ls command 27, 127
 - lsdev -C 48
 - LUN 49, 57
 - LVM 100
- M**
- manual backup 194
 - mapping 9
 - defined 25
 - Maximum jobs per client 97
 - Maximum multiplexing per drive 97
 - maximum pathname length 82, 96
 - Maximum Snapshots (Instant Recovery only) 90, 114, 128, 138, 158
 - MDS 147, 150
 - Media multiplexing 97
 - media server (see NetBackup Media Server)
 - messages file 215
 - method
 - selecting offhost backup 80
 - selecting snapshot 86
 - mirror 6
 - access time 98
 - association 167
 - compared to copy-on-write 7
 - defined 25
 - fast resynch 151
 - overview 6
 - preparing 135
 - rotation 127
 - VxVM snapshot 26, 150
 - mixed environments
 - installing in 34
 - mklogdir script 214
 - mklogdir.bat 215
 - mover.conf file 54, 59, 68, 72
 - 67
 - and AIX 22
 - creating 72
 - DISK 64
 - dr keyword 67
 - END 66
 - examples 73
 - GROUP 65
 - how determine entries for 59
 - hr keyword 64, 67
 - i=reserve value 66
 - order of searching 68
 - policy_name 74
 - TAPE 65
 - multi-device system (VxFS) 147, 150
 - Multiple Copies (Vault) 22
 - multiple data streams 97, 189
 - configuring 97
 - multiplexing 22, 192
 - multi-ported disk arrays 18



N

- name entry in st.conf 51
- NAS
 - offhost backup 9
 - policy configuration 109
 - SnapVault intro 111
- NAS_Snapshot 4, 87, 133, 134, 158, 196
 - access for NetBackup 112
 - backup retention period 158
 - licensing 112
 - logging info 214, 216
 - name 122
 - notes, requirements 112
- NBU_CACHE 136
- nbu_snap method 87, 142, 238
 - with VxVM shared disk group 142
- NDMP 4
 - access web info xvi
 - licensing 112
 - SnapVault intro 111
- NDMP protocol version 112
- NDMP snapshot 158
- ndmp unified log (VxUL) 214, 216
- NDMP V4 133
- NearStore
 - disk type 119
- NetBackup Client Service 113
- NetBackup for NDMP 109
- NetBackup Media Server 9, 25
 - and storage units 81, 192
 - network diagram of 15, 40
 - process overview 19
 - selecting 80
- Network Appliance SnapVault 158
- Network Attached Storage 9, 81
- network interface cards 218
- NEW_STREAM directive 106
- NIC cards and full duplex 218
- no-data Storage Checkpoint 125

O

- Offhost and SAN Data Movement Services
 - (see Advanced Client)
- offhost backup 80
 - and multiplexing 192
 - NAS 9
 - overview 9
 - prerequisites for 192
 - raw partition 192

- SCSI connections 17
 - type of disk (SCSI vs. IDE) 218
 - without SAN 17
- Open Systems SnapVault (OSSV) 116
- operating system
 - changes 51
 - patches 21
- Oracle 21, 230
- OSSV 116
- overwriting
 - raw partition restores 195

P

- P=clientpath 56
- p=devpath 56
- page code 83 22, 56, 58, 192, 218
- pairedisplay listing 178
- pairresync command 224
- partitions
 - Windows 100
- partner host 173
- partner service 174
- passive third-party copy device 40
- passthru device paths 43, 48, 65, 72
- patch for VxFS 21
- patches 21, 217
- pathname length 82, 96
- Perform block level incremental backups 78
- Perform snapshot backups 238
- performance
 - increasing tape 218
- peripherals (latest info on web) xvi
- Persistent Frozen Image (see Instant Recovery)
- PFI_BLI_RESTORE file (for block-level restore) 197
- physical device (as raw partition) 142
- pkgmgr 35
- platform requirements 21
- platforms supported 29
- Point in Time Rollback
 - and SnapVault 121
- point-in-time snapshots 4
- policy
 - how to select type of 80
 - storage unit 81, 97
 - using a wizard to create 97
- Policy dialog 79, 101
- policy_name (on mover.conf file) 68, 72, 74,



- 218
 - poll 171
 - port name
 - using CommandCentral 50
 - primary to mirror association 167
 - primary vs. alternate client
 - processing background 229
 - promotion, file 4, 197, 198
- Q**
- query snapshot 207, 209, 221
 - quiesce 230, 232
- R**
- RAID 5 150, 188
 - RAID Manager daemons 176
 - raw partition 103, 106
 - as snapshot source 96
 - backup 82
 - block vs. character device 192
 - defined 25
 - not supported with VxFS_Checkpoint 147
 - restore 195
 - fsck needed after vxfs restore 196
 - specifying for cache 142
 - recognizing devices 48
 - recovery procedure 221, 238
 - Registry 100, 160
 - and FlashBackup 100
 - rem_drv sg command 49
 - remote snapshot (see alternate client backup)
 - removing
 - clone 226
 - snapshots 221, 238
 - replicated host 14
 - replication
 - for alternate client backup 156
 - testing setup for alternate client backup 157
 - requirements for NetBackup 21
 - reserve/release
 - dr keyword 67
 - hr keyword 64
 - introduction 63
 - third-party reservation 63, 64, 66
 - restore 195
 - and fsck 196
 - block-level restore 197
 - configurations 206
 - FFR with vxvm or FlashSnap 133, 196, 198
 - file promotion 4, 197, 198
 - from disk image 207
 - from FlashBackup 195
 - hardware-level 186, 209
 - logs 215, 216
 - NAS_Snapshot 113
 - Oracle files 197
 - overwrite option 195
 - raw partitions 195
 - re. device file 195
 - Restore everything to its original
 - location 202, 204
 - restrictions 100
 - restrictions (also see the Release Notes) 21
 - resyncfromreplica option 209, 211
 - resynchronization of mirror 151
 - resynchronize
 - at end of backup 186
 - disks 186, 223
 - Resynchronize mirror in background 91
 - retention period 129
 - examples 129
 - proper use 130
 - RMAN 25
 - rollback 196
 - verification 202
 - root, specifying as snapshot source 96
 - rotation 127
 - RSM Database 100, 160
- S**
- s=sn 56
 - SAK 116
 - SAN 22, 192
 - defined 26
 - not required for offhost backup 17
 - SANPoint Control (SPC) 39, 50
 - Schedule tab
 - Instant Recovery 134
 - scripts
 - running before/after backup 92
 - SCSI E4 target descriptors 81
 - SCSI Extended Copy command 54
 - SCSI Inquiry Page Code 83 22, 192, 218
 - SCSI offhost backup 17
 - SCSI reserve/release 63, 64



- dr keyword 67
- hr keyword 64, 67
- i=reserve value 66
- SCSI serialization 22, 192, 218
- SCSI target 49, 58
- SCSI vs. IDE 218
- sd.conf 47
- second world-wide name 57
- serial numbers 22, 56, 192, 218
- serialization 22, 192, 218
- Server Appliance Kit (SAK), Windows 116
- sg.conf 48, 51
- sg.install script 49
- sg.links 48, 51
- sgscan command 43, 48, 62
- Shadow Copy Service (see VSS)
- ShadowImage method 87, 162, 194, 218
- shared disk group (VxVM) 142, 153
- SIGHUP 176
- SIZE_DATA_BUFFERS 218
- Skip verification and force rollback 202, 204
- snap
 - removing 222
- snappcachelist command 241
- snaptcl 100
 - driver 104
 - driver log 215
 - overview 238
- snaplist command 222, 239
- snappoff command 222, 242
- snapon command 239
- snapshot 4
 - auto selection 84, 87, 94, 132
 - back up to local storage 8
 - configuration 86, 131
 - copy-on-write vs mirror (how to choose) 5
 - defined 26
 - deleting 128
 - ID 207, 209, 221
 - instant 136
 - methods, matching to array 163
 - mirror (creating) 153, 155
 - mirror (VxVM) 26, 150
 - mirror access time 98
 - mirror, defined 6
 - naming format
 - NAS snapshot 122
 - on Solaris client, troubleshooting 220
 - overview 4
 - pre-processing for 230
 - removing 221, 238
 - restoring from disk 207, 209
 - rotation 127
 - selecting method of 86
 - source
 - defined 26
 - for symbolic link 96
 - volume (creating) 153, 155
 - VxVM instant 89
- Snapshot Configuration Wizard xvi, 79, 113, 131, 135, 136, 137, 145
- Snapshot Options dialog 84, 86, 94
- Snapshot Policy Configuration Wizard 97
- Snapshot Volume
 - defined 26
- Snapshots only
 - on Schedule tab 132, 134
- snapstat command 242
- SnapVault 20, 114
 - and NDMP protocol 112
 - and OSSV 116
 - and rollback 202
 - and SAK 116
 - changed blocks vs full 120
 - disk type 119
 - enabling NetBackup access 117
 - intro 111
 - logging info 214, 216
 - new subvolume vs existing 120
 - notes 116, 121
 - Open Systems 116
 - prerequisites 116
 - primary and secondary 111
 - primary to secondary access 117
 - restoring 121
 - storage unit 118, 158
 - verification 204
- software
 - distribute to clients 31
 - required (table) 19
 - upgrade 31
- Solaris
 - version requirements 21
- space-optimized snapshots 89, 136, 150, 151
 - snapshot methods supported 152
- SPC 50
- Spectra Logic 62



- split mirror backup (alternate client) 11
- st.conf 47, 51
- stack quiescing 232
- standard (primary) disk 164
- Standard policy
 - restoring 195
- status code 156 217
- Storage Checkpoint 21, 88, 100, 133
 - basic features of 127
 - defined 27
 - determining disk usage 27, 127
 - disk layout 125
- storage devices 78
- storage unit 68, 72, 81, 97
 - for SnapVault 118
 - restrictions for data movers 81, 192
- storage_unit_name version of mover.conf
 - file 75, 218
- streams, allow multiple data 106
- Sun disk array 62
- support web site xvi
- support web site (VERITAS) xvi
- supported
 - data types 100
 - platforms 29
- switched fabric 24
- symbcv command (EMC) 168
- symbolic links 96
- SYMCLI (SYMAPI) 87, 168
- symdbg command (EMC) 168
- symld command (EMC) 168
- Symmetrix disk groups
 - creating 168
- symmir command 169, 186, 223
- synchronize disks 223
- Synchronize mirror before backup 90
- system requirements for NetBackup 21
- System State 100, 160
- system-protected files 160
- system-protected files and FlashBackup 100

T

- tape
 - increasing performance 218
- TAPE keyword 65
- tape reserve 63
- tape-config-list 51
- target (HP limit of 8 devices) 48
- Terminal Services Database 100, 160

- terminate
 - bprd 31
- termination 221, 238
- terminology 23
- third-mirror (vxvm) 150
- third-party copy 9
 - and multiplexing 22, 192
 - and storage unit 81, 97
 - configuration files for 72
 - defined 23
 - device configuration 54, 80
 - network diagram of 16, 41
- third-party copy device
 - active 41
 - and storage units 81, 192
 - defined 27
 - FIXED mode 47
 - how determine passthru path for 59, 62
 - not on same network as media server 42, 69
 - passive 40
 - SCSI reserve 63
- Third-Party Copy Device offhost
 - backup 102
- third-party reservation 64, 66
- TimeFinder method 87, 162, 194, 218
- timeout setting (for mover) 67
- timestamp
 - incremental backup 94
- to keyword 67
- tpautoconf 112
- tpconfig 112
- Transfer throttle 119
- troubleshooting 214
 - directories to create 214
- types of backups supported 102

U

- UFS file system 27, 87, 142
- UID 93
- umount command 208, 223
- UNC for Windows pathnames 110, 115, 117
- unified logging 214, 216
- uninstalling NetBackup 35
- Universal Naming Convention (UNC) 110, 115, 117
- unmount
 - checkpoint 223
 - snapshot 208



- unquiesce 231
- UNSET directive 106, 108
- UNSET_ALL directive 106, 108
- upgrade software 31
- upgrading NetBackup 35
- Use alternate client 80
- user-directed
 - archives 194
 - backup 194

V

- vendors (latest info on) xvi
- verbose mode (btpcinfo) 69
- VERBOSE setting for logs 214
- verification, of rollback 202, 204
- VERITAS Federated Mapping Services 27
- VERITAS Volume Manager 89, 150
- VERITAS Volume Replication 133
- volume
 - defined 27
 - multi-device 147, 150
 - sets (VxVM) 125
- VSP method
 - description 88
 - using with/without Advanced Client 88
- VSS method 88
- VSS_Transportable method 88, 162, 172
 - arrays supported 163
 - environment variables 176
- VVR 21, 125, 133
- VVR method 14, 88, 157
 - preparing for alternate client backup 156
- vxassist 136, 137, 138, 150, 153, 155, 222
- vxassist snapstart 135
- vxdbg command 135, 154, 155, 208, 210
- vxdbg list command 226
- vxdisk command 208
- VxFS clone, removing 222
- VxFS file system 21, 87, 100, 142
 - and AIX, Linux 29
 - patch for library routines 21
 - restoring 196
- VxFS multi-device file system 147, 150
- VxFS_Checkpoint method 88, 100, 133, 147, 197, 198
- VxFS_Snapshot method 89, 104, 149

- vxibc command 157
- vxmake 137
- VxMS 25, 27
- vxprint 136, 137, 222
- vxprint command 222
- vxrecover command 209
- vxsnap 138
- VxVM 188
 - and RAID 5 150
 - clone of disk group 184, 226
 - configuration for arrays 182
 - disk group 182, 186
 - instant snapshots 89, 137, 151
 - mirror 150
 - required version 87, 94, 150
 - restrictions re arrays 182
 - shared disk group 142, 153
 - Volume Manager 21, 150
 - volume name restriction 125, 135, 150
 - volume sets 125
 - volumes (arrays supported) 163
- vxvm method 89, 133, 150, 183, 189, 199
- VxVM mirror
 - preparing for Instant Recovery 135
- vxvol 136, 150, 153

W

- W=wwpn 57
- w=wwpn 57
- web access to recent Advanced Client info xvi
- wildcards in Backup Selections list 82
- Windows
 - OS partitions 100
 - System database files 100, 160
- Windows Server Appliance Kit (SAK) 116
- Windows Shadow Copy Service 159
- wizard
 - Backup Policy Configuration, use of 97
 - for creating a policy 97
- wizard for Advanced Client xvi, 79, 113, 131, 135, 136, 137, 145
- world-wide port name 56, 57

X

- x option (btpcinfo) 42, 69

